

# **The Uni-Level Description: A Uniform Framework for Managing Structural Heterogeneity**

Shawn Bowers

B.S., University of Oregon, 1998

M.S., OGI School of Science and Engineering at OHSU, 2003

A dissertation presented to the faculty of the  
OGI School of Science and Engineering  
at Oregon Health & Science University  
in partial fulfillment of the  
requirements for the degree  
Doctor of Philosophy  
in  
Computer Science and Engineering

January 2004

© Copyright 2004 by Shawn Bowers  
All Rights Reserved

The dissertation “The Uni-Level Description: A Uniform Framework for Managing Structural Heterogeneity” by Shawn Bowers has been examined and approved by the following Examination Committee:

---

Lois Delcambre  
Professor, OGI School of  
Science and Engineering  
Thesis Research Adviser

---

David Maier  
Professor, OGI School of  
Science and Engineering

---

Mark Jones  
Associate Professor, OGI School of  
Science and Engineering

---

Leonidas Fegaras  
Associate Professor, University of  
Texas at Arlington

# Dedication

*To Evening and Elizabeth*

# Acknowledgements

I wish to first thank Lois Delcambre, my advisor, whose help and guidance was invaluable to me throughout my five years of study at the Oregon Graduate Institute (OGI). Without her support, the work in this dissertation would not have been possible. Lois encouraged me to pursue a Ph.D. and introduced me to the subject of databases. She has helped me gain confidence as a researcher and to more effectively communicate my ideas. She also had the difficult task of keeping me on track. It has been a pleasure for me to have the opportunity to work with Lois. And I will always appreciate her patience and the time she has spent as a mentor to me.

Dave Maier also deserves special thanks. He participated in various discussions with Lois and I on issues surrounding meta-data-models and data models. Dave also carefully reviewed many of our papers, always providing detailed and insightful comments. I learned many things from Dave not only in his courses on Scholarship Skills, Object Data Management, and Principles of Database Systems (which Lois and Dave taught together), but also by working with him on papers and proposals, watching him present research, and seeing him discuss and debate ideas in research meetings. It is difficult to articulate the many ways that Lois and Dave have contributed to my personal and professional growth while I attended OGI. I will always be grateful for their help and advice.

My thesis committee also deserves many thanks, which along with Lois and Dave, included Mark Jones and Leo Fegaras. Both Mark and Leo provided valuable, detailed comments and feedback on this dissertation. I also appreciate each members flexibility in terms of scheduling my defense and in receiving drafts of the dissertation while I was in San Diego. In addition, I would like to thank Juliana Freire, Peter Heeman, and Renée Miller for participating in my thesis proposal.

There are many students at OGI, and particularly in the database group, that deserve thanks. I collaborated most closely with Mathew Weaver and Sudarshan Murthy, researching various aspects of superimposed information. Not only did I enjoy working with them, but I was exposed to many aspects of their research, which strengthened my knowledge and understanding in the area. I would also like to thank Vassilis Papadimos and William Howe for being genuinely interested in discussing new research ideas and for engaging in lively debates. Teaching the database-system course with Peter Tucker enriched my education at OGI, and our technical discussions concerning all aspects of database management helped me greatly. I also want to thank the students and faculty that participated in the database reading group for broadening my knowledge of research issues in data management and computer science in general.

Most importantly, I would like to thank my wife Evening who is my best friend and who always believed I could finish. And to my daughter Elizabeth, who could not wait until I was “done working.” The rest of my family also deserves thanks: my grandfather for gracefully supporting my education and showing endless faith that I can accomplish whatever I set out to do, my grandmother for her support and understanding, and special thanks to my parents for their generous help and encouragement.

During the course of this work, I was supported in part through research grants provided by the National Science Foundation (grant numbers EIA-9983518 and IIS-9817492) and the Defense Advanced Research Projects Agency (contracts F30602-96-1-0302 and N66001-00-C-8032). Finally, I would like to thank Bertram Ludäscher, who graciously helped make it possible for me to finish writing this dissertation and to schedule my thesis defense while I was working for him at the San Diego Supercomputer Center.

# Contents

|                                                              |             |
|--------------------------------------------------------------|-------------|
| <b>Dedication</b> . . . . .                                  | <b>iv</b>   |
| <b>Acknowledgements</b> . . . . .                            | <b>v</b>    |
| <b>Abstract</b> . . . . .                                    | <b>xiii</b> |
| <b>1 Introduction</b> . . . . .                              | <b>1</b>    |
| 1.1 Structural Heterogeneity . . . . .                       | 2           |
| 1.2 Managing Structural Heterogeneity . . . . .              | 6           |
| 1.3 Our Approach . . . . .                                   | 8           |
| 1.4 Outline . . . . .                                        | 10          |
| <b>2 The Uni-Level Description</b> . . . . .                 | <b>12</b>   |
| 2.1 Architecture . . . . .                                   | 13          |
| 2.1.1 Typical Meta-Data-Model Architectures . . . . .        | 13          |
| 2.1.2 The ULD Architecture . . . . .                         | 15          |
| 2.2 The ULD Meta-Data-Model . . . . .                        | 17          |
| 2.3 The ULD Representation Language . . . . .                | 18          |
| 2.4 The ULD Query Language . . . . .                         | 27          |
| 2.4.1 Datalog . . . . .                                      | 30          |
| 2.4.2 Querying ULD Configurations using Datalog . . . . .    | 32          |
| 2.5 The ULD Constraint Language . . . . .                    | 35          |
| 2.6 A Partial ULD Axiomatization . . . . .                   | 42          |
| 2.6.1 Definition of the Core ULD Theory . . . . .            | 44          |
| 2.6.2 Definition of the ULD Meta-Data-Model Theory . . . . . | 47          |
| 2.7 Summary . . . . .                                        | 52          |
| <b>3 Modeling Data Models</b> . . . . .                      | <b>54</b>   |
| 3.1 Expressing Common Constraints . . . . .                  | 55          |
| 3.1.1 Uniqueness Constraints . . . . .                       | 55          |
| 3.1.2 Collection Constraints . . . . .                       | 57          |

|          |                                                            |            |
|----------|------------------------------------------------------------|------------|
| 3.1.3    | Inclusion Constraints . . . . .                            | 59         |
| 3.1.4    | Digression: Regular Subset Inclusion Constraints . . . . . | 63         |
| 3.2      | Traditional Database Data Models . . . . .                 | 64         |
| 3.2.1    | Relational Data Models . . . . .                           | 65         |
| 3.2.2    | Object-Oriented Data Models . . . . .                      | 69         |
| 3.2.3    | Entity-Relationship Data Models . . . . .                  | 73         |
| 3.3      | Structured Data Models for the Web . . . . .               | 79         |
| 3.3.1    | XML Data Models . . . . .                                  | 79         |
| 3.3.2    | RDF Data Models . . . . .                                  | 88         |
| 3.4      | Additional Data Models . . . . .                           | 93         |
| 3.4.1    | Topic-Map Data Models . . . . .                            | 94         |
| 3.4.2    | The VIKI Data Model . . . . .                              | 97         |
| 3.5      | Properties of Conformance . . . . .                        | 101        |
| 3.5.1    | Structural Conformance Properties . . . . .                | 101        |
| 3.5.2    | Process-Oriented Conformance Properties . . . . .          | 105        |
| 3.5.3    | The ULD Conformance Properties . . . . .                   | 105        |
| 3.6      | Summary . . . . .                                          | 106        |
| <b>4</b> | <b>Transformation . . . . .</b>                            | <b>109</b> |
| 4.1      | The ULD Transformation Language . . . . .                  | 109        |
| 4.2      | Transformation Examples . . . . .                          | 114        |
| 4.2.1    | E-R to Relational Transformations . . . . .                | 115        |
| 4.2.2    | XML to Relational Transformations . . . . .                | 118        |
| 4.2.3    | Schema Transformation . . . . .                            | 125        |
| 4.2.4    | Transformations for Standard Serializations . . . . .      | 127        |
| 4.3      | Analyzing ULD Transformations . . . . .                    | 131        |
| 4.3.1    | Well-Typed Transformations . . . . .                       | 131        |
| 4.3.2    | Mapping Rules and Data-Model Constraints . . . . .         | 138        |
| 4.3.3    | Cardinality Constraints on Conformance . . . . .           | 140        |
| 4.3.4    | Inference and Transformation . . . . .                     | 141        |
| 4.4      | Summary . . . . .                                          | 145        |
| <b>5</b> | <b>Incremental Navigation . . . . .</b>                    | <b>148</b> |
| 5.1      | Motivating Examples . . . . .                              | 149        |
| 5.2      | Navigation Operators . . . . .                             | 150        |
| 5.3      | Navigation Bindings as ULD Queries . . . . .               | 153        |
| 5.4      | Higher-Level Binding Specifications . . . . .              | 155        |

|                           |             |
|---------------------------|-------------|
| <b>6 Conclusion</b>       | <b>.161</b> |
| 6.1 Implementation Issues | 162         |
| 6.2 Related Work          | 166         |
| 6.3 Contributions         | 169         |
| 6.4 Future Work           | 171         |
| <b>Bibliography</b>       | <b>.174</b> |
| <b>Biographical Note</b>  | <b>.182</b> |

# List of Tables

|     |                                                                 |     |
|-----|-----------------------------------------------------------------|-----|
| 3.1 | Structural conformance properties of data models. . . . .       | 102 |
| 3.2 | Process-oriented conformance properties of data models. . . . . | 106 |

# List of Figures

|      |                                                                                                                      |    |
|------|----------------------------------------------------------------------------------------------------------------------|----|
| 1.1  | An example of a relational schema and instance. . . . .                                                              | 2  |
| 1.2  | An example of an XML schema and instance. . . . .                                                                    | 3  |
| 1.3  | An example of an RDF schema and instance. . . . .                                                                    | 4  |
| 1.4  | The ULD-enabled environment where generic tools can access information<br>stored in various data models. . . . .     | 10 |
| 2.1  | A standard meta-data-model architecture. . . . .                                                                     | 13 |
| 2.2  | The ULD meta-data-model architecture. . . . .                                                                        | 16 |
| 2.3  | The standard meta-data-model architecture (left) captured in the ULD ar-<br>chitecture (right). . . . .              | 16 |
| 2.4  | An example of data without corresponding schema represented in the ULD.                                              | 17 |
| 2.5  | The syntax for ULD construct and instance expressions. . . . .                                                       | 19 |
| 2.6  | The default <code>uldValue</code> and <code>uldValuetype</code> constructs for representing scalar<br>types. . . . . | 24 |
| 2.7  | The interpretation of ULD configurations. . . . .                                                                    | 35 |
| 2.8  | The naming conventions used in specifying ULD axioms. . . . .                                                        | 44 |
| 2.9  | The core ULD axioms for describing basic structure restrictions. . . . .                                             | 45 |
| 2.10 | The core ULD axioms describing conformance cardinality restrictions. . . .                                           | 46 |
| 2.11 | The ULD meta-data-model axioms that define the allowable construct types<br>and two structural restrictions. . . . . | 47 |
| 2.12 | The axioms for tuple structures in the ULD meta-data-model. . . . .                                                  | 48 |
| 2.13 | Two general axioms for collection structures in the ULD meta-data-model.                                             | 49 |
| 2.14 | The axioms for set structures in the ULD meta-data-model. . . . .                                                    | 49 |
| 2.15 | The axioms for bag structures in the ULD meta-data-model. . . . .                                                    | 50 |
| 2.16 | The axioms for list structures in the ULD meta-data-model. . . . .                                                   | 51 |
| 2.17 | The axioms for union definitions in the ULD meta-data-model. . . . .                                                 | 52 |
| 3.1  | An example of a Topic Map represented graphically. . . . .                                                           | 95 |
| 3.2  | An example of a hypertext and its type definition in VIKI. . . . .                                                   | 99 |

|     |                                                                                                                                                                                         |     |
|-----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 4.1 | The edge-based, tree representation of the XML document of Figure 1.2. Nodes are annotated with their ordinal number and levels are annotated with their <i>sourceId</i> value. . . . . | 119 |
| 4.2 | The result of applying the edge transformation. . . . .                                                                                                                                 | 120 |
| 4.3 | The result of applying the attribute transformation. . . . .                                                                                                                            | 123 |
| 4.4 | The result of applying the RDF serialization mapping. . . . .                                                                                                                           | 128 |
| 4.5 | The basic type annotation rules for mappings. . . . .                                                                                                                                   | 134 |
| 4.6 | The collection type annotation rules for mappings. . . . .                                                                                                                              | 134 |
| 4.7 | The selector, conformance-cardinality, and identifier-generation type annotation rules for mappings. . . . .                                                                            | 135 |
| 4.8 | The typing rules for ULD mappings. . . . .                                                                                                                                              | 137 |
| 5.1 | A high-level binding for direct navigation of RDF. . . . .                                                                                                                              | 155 |
| 5.2 | The Datalog rules to compute links and locations based on a binding-constructs specification. . . . .                                                                                   | 158 |
| 5.3 | An example of labeled instances for RDF. . . . .                                                                                                                                        | 159 |
| 6.1 | The basic architecture of the ULD transformation engine. . . . .                                                                                                                        | 163 |
| 6.2 | The basic architecture of the ULD navigation prototype. . . . .                                                                                                                         | 165 |
| 6.3 | An architecture for ULD transformation using non-materialized views. . . .                                                                                                              | 166 |

# Abstract

## **The Uni-Level Description: A Uniform Framework for Managing Structural Heterogeneity**

**Shawn Bowers**

**Supervising Professor: Lois Delcambre**

Information management systems (such as database and knowledge-based systems) are based on data models, which provide basic structures for organizing and storing data. A number of data models are in use, and each provides a slightly different set of structures. For instance, information is represented as tables in the relational data model, as ordered trees in semi-structured data models (such as XML), and as directed graphs in semantic networks (such as RDF). With several distinct data models available, developers can select the most convenient representation for their particular need. However, the use of multiple data models introduces structural heterogeneity, making it difficult to combine information and exploit generic tools (for example, for querying or browsing).

This dissertation describes a new framework called the Uni-Level Description (ULD) that can accurately store and accommodate information from a broad range of data models. The ULD consists of a meta-data-model to describe the basic data structures used by a data model, and a uniform representation to store information within a source, including the data-model structures, schema (if present), and instance data. The ULD extends existing meta-data-models by allowing uniform access to schema and data, by permitting data models with non-traditional schema arrangements, and by providing a language for representing data-model constraints.

The two primary motivations for our work are (1) to study the basic structural capabilities offered by data models and (2) to define languages for enabling interoperability among sources with structural heterogeneity. We use the ULD to describe a wide range of data models, including those that allow optional and partial schema. We present a query and transformation language (based on Datalog) for accessing and converting information among heterogeneous sources. We demonstrate the flexibility of the transformation language by defining a number of structural mappings. And finally, we use the ULD as the basis for generic navigation, where a single set of operators can be used to discover and browse information in structurally heterogeneous sources.

# Chapter 1

## Introduction

Information management systems provide a range of tools for working with data. For example, most systems provide tools to create, manipulate, and query data, and to maintain data consistency. Although not always obvious, the most fundamental tool an information system provides is its *data model*—a fixed set of data structures for storing, organizing, and processing information.

A data model defines a language for representing both conceptual information (in a *schema*) and instance data. A schema describes a domain of interest and typically prescribes an organization for data. Schema and data are represented in a data model using named data structures, called *constructs* [79]. The challenge when creating a data model is to select an appropriate mix of constructs that can both conveniently represent a wide range of information and enable tools for efficiently managing data.

In this dissertation, we study the structuring capabilities of data models. A large number of data models exist and are being used, and each provides a different set of constructs for representing information. Part of our motivation for examining data models is to enable interoperability—we envision a generic framework that can support tools for exploiting information, regardless of the data model used for representation. Here, we focus on the structural differences among data models and provide languages for expressing and resolving these differences.

This chapter introduces the problem of structural heterogeneity, which arises in part from the use of different data models, describes existing approaches for managing structural heterogeneity, introduces our approach (which we call the *Uni-Level Description*), and concludes with an outline of the rest of this dissertation.

*Movie*

| <u>mid</u> : <i>integer</i> | title : <i>string</i> | runtime : <i>integer</i> | genre : <i>string</i> | company : <i>string</i> |
|-----------------------------|-----------------------|--------------------------|-----------------------|-------------------------|
| 1                           | 'The Usual Suspects'  | 106                      | 'Thriller'            | 'Gramercy'              |
| 2                           | 'Meet the Parents'    | 108                      | 'Comedy'              | 'Universal'             |

*Cast*

| <u>mid</u> : <i>integer</i> | character : <i>string</i> | actor : <i>string</i> |
|-----------------------------|---------------------------|-----------------------|
| 1                           | 'Roger Verbal Kint'       | 'Kevin Spacey'        |
| 2                           | 'Jack Byrnes'             | 'Robert De Niro'      |

*Review*

| <u>mid</u> : <i>integer</i> | rating : <i>float</i> | source : <i>string</i> |
|-----------------------------|-----------------------|------------------------|
| 1                           | 8.7                   | 'IMDB'                 |
| 2                           | 7.0                   | 'IMDB'                 |

Figure 1.1: An example of a relational schema and instance.

## 1.1 Structural Heterogeneity

Figures 1.1, 1.2, and 1.3 show similar schema and instance data represented in the relational [2, 33, 71], XML [22], and RDF [47] data models. We describe these data models below to highlight the basic structural differences that occur when similar information is represented in disparate data models.

One of the most widely used data models is the relational model. The primary construct of the relational model is a table, which is similar to a mathematical relation. As shown in Figure 1.1, a relational schema consists of a set of table definitions, typically called *relations*, containing attribute names and corresponding scalar types, called *domains*. A table instantiates a relation, and consists of tuples, where each tuple has exactly one value for each attribute. Figure 1.1 shows three table definitions and their corresponding tables. A table definition can designate one or more attributes as a primary key, where every tuple in the table has a distinct collection of values for the attributes. Primary keys are usually underlined, as shown in Figure 1.1. A foreign key may also be defined between attributes in two relations, where a set of values in one table contains key values in the other table, forming a value-based relationship between the tables. (Foreign keys could be defined for Figure 1.1, for example, between the Cast and Movie, and Review and Movie tables.)

The Extensible Markup Language (XML) is a more recent data model that is used for

*DTD Schema*

```

<!ELEMENT moviedb (movie*)>
<!ELEMENT movie (title,studio,genre*,actor*)>
<!ELEMENT title (#PCDATA)>
<!ELEMENT studio (#PCDATA)>
<!ELEMENT genre (#PCDATA)>
<!ELEMENT actor (#PCDATA)>
<!ATTLIST actor role CDATA #REQUIRED>

```

*XML Document*

```

<?xml version="1.0"?>
<moviedb>
  <movie>
    <title>The Usual Suspects</title>
    <studio>Gramercy</studio>
    <genre>Thriller</genre>
    <actor role="Supporting Actor">Spacey, Kevin</actor>
  </movie>
  <movie>
    <title>Meet the Parents</title>
    <studio>Universal</studio>
    <genre>Comedy</genre>
    <actor role="Leading Actor">De Niro, Robert</actor>
  </movie>
</moviedb>

```

Figure 1.2: An example of an XML schema and instance.

information interchange. In XML, a document is represented as a text file that defines a labeled, ordered tree. Schema languages for XML are based on tree grammars, which specify regular expressions over node labels. An example schema for XML (expressed as a document type definition (DTD) [22]) and an instance document are shown in Figure 1.2. Schemas for XML consist of element and attribute definitions. An element represents a node in a tree that can contain a set of attribute-value pairs and has an ordered collection of children nodes, where each child node is either an element or character-string data.

The Resource Description Framework (RDF) is the primary representation scheme for the Semantic Web [11]. The RDF data model is based on directed graphs, similar to semantic networks [53, 82], and is designed for annotating Web-content with semantic information. An example schema and instance are shown in Figure 1.3. Nodes in the graph are called *resources*, which are connected using *properties*. In RDF, statements about




Figure 1.3: An example of an RDF schema and instance.

resources are made using *triples*, which relate a subject resource to an object resource through a property (sometimes called a *predicate* in a triple). RDF Schema (RDFS) [23] is a schema language for RDF that provides constructs for defining classes, subclasses, properties, and subproperties. As shown in Figure 1.3, the schema information is located at the top of the figure, where subclasses are drawn using a special relationship (for example, the “oscar” class is a subclass of “award”). Instances are located at the bottom of the figure, where relationships labeled with `rdf:type` connect instances to classes.

The relational model, XML, and RDF are structurally different data models. The relational model represents information using relations, XML using trees, and RDF using graphs.<sup>1</sup> One advantage of having several different data models is that users can select the right representation for their particular need, for example, based on the structural properties of the data they wish to represent. However, multiple representations introduce structural, model-based heterogeneity, making it difficult to combine information from different sources or to exploit information using generic tools, for example, for querying or browsing.

<sup>1</sup>Although not discussed in the previous examples, these data models also have fundamental differences in the way schemas restrict instance data, which we describe further throughout this dissertation.

Structural heterogeneity [40] can occur at various levels within a database. The presence of different data models introduces fundamental differences in how similar information is structured. Different schemas can also exhibit structural heterogeneity. Namely, two information sources may contain data from similar domains, but differ in how each schema organizes and classifies information. For example, the schemas for XML and RDF shown in Figures 1.2 and 1.3, respectively, have the following differences, even though they describe the same basic domain.

- Each schema includes information not found in the other. For example, the XML schema defines studio information (that is, the company that created the movie), whereas the RDF schema defines award information (that is, the awards won by the movie).
- The schemas represent conceptually identical information using different names and structural types. For example, the XML schema uses the label “movie,” whereas the RDF schema uses “film.” In addition, the RDF schema defines a film title as a direct property of a film and a film is allowed to have multiple titles. In the XML schema, a title is modeled as a sub-element of a movie and a movie has exactly one title. Note also that, in the RDF example, film and title are represented using different data-model constructs, whereas in the XML example, the same data-model construct is used to represent title and movie.
- The schemas contain *schematic* differences [57], where schema information in RDF is represented as data in XML. In particular, the RDF schema represents a genre as a class in the schema, whereas in XML a genre is represented as a data item (in this case, as plain text).

Finally, individual data values can also exhibit structural heterogeneity. For example, in the XML instance of Figure 1.2, the names of actors are stored last-name first, whereas in the RDF instance, the names of actors are stored first-name first. This example is a simple case, and in general, differences in data values can be complex.

## 1.2 Managing Structural Heterogeneity

Most approaches for managing structural heterogeneity focus on resolving differences in schema as opposed to differences in data models. These approaches, which we discuss here, resolve heterogeneity by providing support for either *schema transformation* or *schema integration*.

Schema transformation approaches [1, 12, 32, 35, 68, 69] provide techniques to convert valid instance data under one schema to valid instance data under another schema, where the source and target schemas are assumed to be heterogeneous. These approaches typically define a special-purpose, query-based transformation language to express schema mappings that, when executed, perform the desired instance-data conversion. In addition, a number of approaches attempt to determine schema mappings [70] automatically (or semi-automatically) by finding matching names and structures within the schemas.

Schema integration approaches [10, 24, 38, 40, 52, 62, 64, 75] focus on combining information sources with heterogeneous schemas into a single collection of data with a single schema (called a *global* schema). Many approaches [24, 52, 62, 75] construct a virtual integrated collection, where a query against the global schema is executed by a *mediator*, which rewrites the query into one or more local queries (using the local schemas), combines the results of each local query, and returns the integrated result. Schema integration requires schema transformation to resolve heterogeneity. In particular, schema mappings are defined as views (expressed as queries) between local and global schemas.

Approaches for schema transformation and integration typically assume all schemas and (possibly) data are represented in, or converted to the same, common data model. The common data model is often designed specifically for the approach.

In general, however, developers still use *ad hoc* solutions for resolving structural heterogeneity [12]. By *ad hoc*, we mean special-purpose programs, typically written in procedural languages, for resolving specific instances of structural heterogeneity. Even in schema transformation and integration approaches, special-purpose programs are used to map between data models, for example, to convert information in and out of a common data model. Writing programs to convert information between data models is often a

time-consuming and error-prone task that leads to complex programs that are difficult to maintain and reuse. The following list highlights some of the reasons why converting information between data models, especially using *ad hoc* solutions, is difficult. We use the term *data-model mapping* to denote a mapping from information represented in one data model to information represented in another data model, which in general may also require schema mapping.

- *Data models contain different modeling constructs.* Although obvious, having disparate constructs means that mappings between data models are rarely one-to-one, and there are usually many different and equally “correct” ways to convert information between data models. For example, numerous approaches have been proposed for converting XML into the relational model [14, 77], where each conversion has benefit in certain situations.
- *There may be different versions of the same data model.* For example, many versions of the E-R data model have been defined [27, 28], each with slightly different features. In particular, some E-R models permit only binary relationships, whereas others allow  $n$ -ary relationships; and some E-R data models permit both entities and relationships to have attributes, whereas others do not allow attributes on relationships.
- *Desired mappings between data models may be partial.* In particular, some constructs or schema types may not be applicable to the target information source. In general, partial mappings result in “lossy” transformations, where the original database cannot be rebuilt from the converted information.
- *Mappings between data models may require encoding conventions.* Because data-model constructs differ, *conventions* are often used to represent the desired information within a target data model, when the target data model does not contain the necessary constructs. For example, conventions are heavily used in most XML-to-relational mappings, because the relational model does not have explicit constructs for representing trees. The use of conventions introduces implicit assumptions and

dependencies between source and target databases.

- *Not all data models require schema, and some allow multiple levels of schema.* In general, data models that allow non-traditional schema arrangements are problematic for schema transformation and integration approaches, because not all data of interest may have corresponding schema information, and schemas may be further organized by meta-schema structures [43, 59]. By non-traditional, we mean schema arrangements not typically allowed in traditional data models, such as permitting partial, optional, and multiple levels of schema. Similarly, schematic differences and differences in how data is represented also generally require complex transformations [3, 57].
- *Data models have inherent constraints.* Resolving structural heterogeneity is complicated by data-model constraints. For example, *ad hoc* solutions must be aware of constraints to avoid incorrect conversions. Data-model constraints can range from simple, for example, relational tables have a corresponding table definition, to complex, for example, the restriction placed on tables under primary and foreign keys.

### 1.3 Our Approach

In this dissertation, we describe a high-level, declarative language for expressing and manipulating data models, called the *Uni-Level Description* (ULD) [19, 20]. With the ULD, it is possible to describe the constructs and constraints of a wide range of data models. In addition, the ULD provides uniform access to data-model constructs, schema, and instance-data, which enables flexible transformation rules for resolving structural heterogeneity and provides a means to implement generic operations over data models.

The ULD is motivated by a simple observation: most data models use only a small set of basic data structures, such as *atomic*, *tuple*, and *collection* types, composed in various ways, to store information. Our approach is to describe the constructs of a data model using these basic structures and then instantiate them to describe the schema and instance information that is present in a data source. The goals of the ULD are as follows.

- Permit “faithful” representations of a wide variety of data models. For example, the relational model should be described as having relations, tables, tuples, and so on, without requiring these constructs to be converted to the constructs of another, common data model.
- Allow a single and complete representation of an information source, including its data, schema(s) if present, and data model.
- Provide uniform access to information within a source. In particular, the ULD should be a uniform representation, much like RDF, so that query or navigation operators can access the data model and schema information as easily as data. More than that, the data model, schema, and data information should be accessible at the same time, for example, in the same transformation rule.
- Accommodate information with flexibility regarding the use of schemas. We want to represent information where the schema is required (as in the relational model), where schema can be missing (as in XML), and where there can be multiple levels of schema (as in RDF).
- Support data-model constraints. The ULD should provide an expressive enough language to capture inherent data-model constraints.

A key characteristic of the ULD is its use of explicit and distinct *instance-of* relationships for representing information. These relationships enable uniform access to information while clearly distinguishing between data-model constructs, schema (when appropriate), and data. This approach differs from data models that overload the *instance-of* (or *type*) relationship to provide uniform representation (such as RDF, Telos [43, 59], or other semantic-network-based data models). In addition, having distinct instance-of relationships allows a data-model configurer to place explicit constraints on schema-instance conformance, permitting accurate descriptions of a wide range of data models.

We envision an environment where arbitrary data sources are easily described using the ULD, enabling the use of generic, ULD-based tools that can work over arbitrary information sources, as shown in Figure 1.4. We describe two such tools that exploit the ULD




Figure 1.4: The ULD-enabled environment where generic tools can access information stored in various data models.

in this way. The first is a transformation language that allows powerful transformation rules expressed in Datalog (or Prolog) for defining and executing mappings to convert information from one source to another in the presence of different schemas or data models. The other is a browser, supported by a navigational API against the ULD, that permits both naïve users (for example, users who do not know, or need to know, what an XML element or attribute is) and sophisticated users (for example, agents or crawlers) to access information sources in a generic manner.

## 1.4 Outline

The rest of this dissertation is organized as follows. In Chapter 2 we describe the ULD and the associated languages for representing information sources, expressing queries (to retrieve information represented in the ULD), and specifying constraints. In Chapter 3 we demonstrate the ULD by giving complete descriptions for a wide range of data models. Chapter 3 also provides properties to classify data models in terms of the types of schema arrangements they allow. The classification is inspired by the ULD architecture and its flexibility in describing schema-instance relationships. In Chapter 4 we introduce the ULD transformation language, which is an application of the ULD architecture. We demonstrate the ability of the transformation language to express flexible mappings for

resolving structural heterogeneity, including differences in data models and schemas. In Chapter 5 we describe a navigation interface for generically browsing information represented in disparate data models. The navigation interface is an example of a generic tool: the same operators are used to navigate information regardless of the data model used for representation. Finally, in Chapter 6 we conclude by discussing ULD implementation issues, related work, our contributions, and future work.

# Chapter 2

## The Uni-Level Description

In this chapter, we introduce the *Uni-Level Description* (ULD), which is both a *meta-data-model* (that is, it can describe data models), and a distinct representation scheme: it can directly represent both schema and instance information (expressed in terms of data model constructs). We also describe the ULD query and constraint languages, and give a formal definition of the underlying theory of the ULD.

This chapter is organized as follows. In Section 2.1, we describe existing meta-data-model architectures and contrast them with the architecture of the ULD. In Section 2.2, we present the ULD meta-data-model (that is, a specific set of basic structures for the ULD), which we use in the rest of this dissertation. In Section 2.3, we describe an implementation-independent syntax for representing information in the ULD and use it to describe the relational, XML, and RDF data models. In Section 2.4, we describe a declarative query language for the ULD based on Datalog. The ULD query language provides a convenient, high-level language for uniformly accessing data model, schema, and instance data. We also give a formal, logic-based interpretation of the ULD representation language. In Section 2.5, we describe a constraint language for the ULD, which extends the ULD query language of Section 2.4, and can be used to specify data-model constraints. Finally, in Section 2.6, we give an axiomatization for the ULD using the constraint language, which is enabled by the flexibility of the ULD architecture, and the ability of the query language to uniformly access constructs, schema, and data in the ULD.




Figure 2.1: A standard meta-data-model architecture.

## 2.1 Architecture

### 2.1.1 Typical Meta-Data-Model Architectures

The standard architecture for representing information consists of three levels: data model, schema, and (instance) data. A *meta-data-model* adds a new level of abstraction to the top of this architecture to describe one or possibly more data models, as shown in Figure 2.1. (Note that the architecture of Figure 2.1 is more general than most existing architectures [5, 6, 8, 31, 32, 78].) As shown, the architecture consists of four levels. The top level represents the meta-data-model, the second level from the top represents the meta-schema (that is, the data model), the third level represents schema, and the bottom level represents instance data. A meta-schema defines the set of structural *constructs* provided by a data model for representing schema; and schemas are represented as instantiations of the meta-schema constructs. For example, the meta-schema for the relational data model would describe relations, attributes, primary and foreign keys, and so on.

In a *self-describing data model* [54], the meta-data-model and the data model are the same—that is, the same storage structures are used for schema and data. For example, a self-describing representation of the relational data model would encode a relational schema (that is, the set of relation names, their attributes, foreign and primary keys, and domains) using a fixed set of relations, called the *catalog*. The catalog can then be queried, for example, to find the attributes of a relation.

In contrast to a self-describing data model, a meta-data-model can be used to enable

interoperability among heterogeneous sources that use disparate data models (as opposed to just the relational model, for example). In this case, each heterogeneous source is described using the same meta-data-model, which provides a common representation for accessing schema (via the meta-schema structures).

A number of approaches use E-R (entity and relationship) structures to define meta-schemas and their corresponding schemas [5, 6, 8, 31, 61, 78]. In this case, meta-schema constructs are defined as patterns, or compositions of entity and relationship types. These types are used to define entities and relationships representing corresponding schema items. Using this approach, a meta-schema for an object-oriented data model would contain a class construct represented as an entity type (called *class*) with a name attribute and relationships to other constructs (for example, that represent class attributes). A particular object-oriented schema is then represented as a set of entities and relationships that instantiate these meta-schema types. We note that, in most meta-data-model architectures (YAT [32] being an exception), data (the bottom level of Figure 2.1) is stored in the database system, and is not represented in the meta-data-model architecture explicitly.

An important assumption made by existing meta-data-model approaches is that data (the bottom level of Figure 2.1) must adhere to the schema (the third level of Figure 2.1). That is, all data is considered to be an instance of existing schema definitions. This assumption stems from traditional database systems, which generally require *complete schema*. That is, in a traditional database system, all data must satisfy all of the constraints imposed by a schema. However, a number of data models also permit *optional* and *partial* schema definitions, in which some or all of the data does not conform to a defined schema and, in the case of a partial schema, data items that conform to a schema can be mixed with data items that do not have associated schema definitions. Optional- and partial-schema data models cannot be represented with the assumptions underlying the architecture of Figure 2.1. Namely, there is no way to go from meta-schema directly to data without first going through schema—the meta-schema only defines structures for representing schema.

### 2.1.2 The ULD Architecture

In contrast to Figure 2.1, the ULD uses the architecture<sup>1</sup> shown in Figure 2.2. The ULD meta-data-model (shown as the top level in Figure 2.2) consists of *construct types* (that is, meta-constructs) that denote structural primitives. The middle level uses the structural primitives to define both data and schema *constructs*, possibly with *conformance relationships* between them. For example, in Figure 2.2, the *class* and *object* constructs are related through a conformance link (labeled *conf*) and their corresponding *construct instances* in the bottom layer (that is, *actor* and the object with the name ‘Robert De Niro’) are related through a data instance-of link (labeled *d-inst*). It is important to note that both schema and data constructs of a data model are explicitly represented in the ULD. In existing meta-data-model architectures, only schema constructs (as shown in Figure 2.1) are defined in the meta-schema—data constructs are assumed to implicitly follow from the definition of the meta-schema. The bottom level of Figure 2.2 represents schema and data as instances of constructs, connected by *d-inst* links when appropriate. The relationship between the ULD architecture and the architecture of Figure 2.1 is shown in Figure 2.3. In particular, the ULD architecture subsumes the typical meta-data-model approaches of Figure 2.1.

The ULD distinguishes three kinds of instance-of relationships. Constructs introduced in the middle layer are necessarily instances of construct types in the meta-data-model, represented with *ct-inst* instance-of links. Similarly, every item introduced in the bottom layer is necessarily an instance of a construct in the middle layer, represented with *c-inst* instance-of links. Finally, an item in the bottom layer can be an instance of another item in the bottom layer, represented with *d-inst* instance-of links, as allowed by *conformance* relationships specified in the middle layer. Note that we distinguish the term *instance* from the term *instance-of*. As an example, the bottom layer of the ULD architecture

---

<sup>1</sup>The ULD is a flat representation in that all information stored in the ULD is uniformly accessible as a single level. For example, a single query or transformation rule can access data-model constructs, schema, and instance data. The name *Uni-Level Description* was chosen to reflect this property of the ULD, that is, of being a flat representation. However, as shown in Figure 2.2, information stored in the ULD is *logically* divided into three layers using explicit instance-of relationships (which we describe in this chapter).




Figure 2.2: The ULD meta-data-model architecture.




Figure 2.3: The standard meta-data-model architecture (left) captured in the ULD architecture (right).

consists of *construct instances* (that is, data and schema items), and items within the architecture are associated using *instance-of* relationships (that is, *ct-inst*, *c-inst*, and *d-inst*). A construct instance may be a *d-inst* of another construct instance (all in the bottom layer), and must be a *c-inst* of at least one construct (in the middle layer).

The ULD is able to represent a wide range of data models using the flexibility offered by the *conf* and *d-inst* relationships. For example, in XML, elements are not required to conform to element types. In this case, the XML element would not have a *d-inst* link to an XML element type (because there are no element types explicitly defined that the element conforms to). Figure 2.4 shows the case where data is represented in the ULD without corresponding schema.

The ULD represents an information source as a *configuration*, which contains the




Figure 2.4: An example of data without corresponding schema represented in the ULD.

constructs of a data model, the construct instances (both schema and data) of a source, and the associated conformance and instance-of relationships. Thus, a configuration can be viewed as an instantiation of Figure 2.2. Each configuration uses a finite set of identifiers to denote construct types, constructs, and construct instances as well as a finite set of *ct-inst*, *c-inst*, *conf*, and *d-inst* relationships. Together, the use of identifiers with explicit instance-of relationships allows all three levels of information in Figure 2.2 to be stored in a single configuration—enabling direct and uniform access to all data-model, schema, and data information in a source. In a ULD configuration, all construct and construct-instance identifiers have exactly one associated value, where a value instantiates a meta-data-model structure.

## 2.2 The ULD Meta-Data-Model

This section briefly describes the choice of structures we use as the construct types of the ULD meta-data-model. We describe these structures in more detail in the following section.

The ULD meta-data-model contains primitive structures for tuples, that is, sets of name-value pairs; set, list and bag collections; atomics, for scalar values such as strings and integers; and unions, for representing non-structural, generalization relationships among constructs. The construct-type identifiers for these structures are denoted *struct-ct*, *set-ct*, *list-ct*, *bag-ct*, *atomic-ct*, and *union-ct*, respectively.

Scalar values such as strings and integers are treated as subsets of the construct-instance identifiers in a configuration. This choice of representation leads to more concise configurations, however, primitive values could also be represented as distinct values, disjoint from the set of identifiers.

Tuple and collection structures can be nested, but only through identifiers. For instance, set values always contain a unique collection of identifiers, as opposed to a unique collection of values. Thus, to construct a nested set value, we would create an identifier  $i_1$  whose value is a set containing an identifier  $i_2$ , where  $i_2$ 's value is a set.

The structures of the ULD meta-data-model are similar to those found in E-R [27] and object-oriented database (OODB) [48] data models. Unlike the E-R data model, however, the ULD does not treat relationships as first class; a relationship in the ULD is specified within a tuple structure, which is similar to some OODB data models. The ULD also has explicit structures for representing collections, which have identifiers and can stand alone or be nested within tuples and other collections. Thus, collections in the ULD serve to represent extents, set-valued relationships, and collection-valued attributes in OODB data models. Finally, the ULD uses union types, and does not provide *isa* relationships between structures, which are found in OODB and some E-R data models. We note that other meta-data-models, for example, with different structures or allowable compositions of structure, are possible without requiring changes to the ULD architecture of Figure 2.2. However, the particular meta-data-model we have chosen is rich enough to express a wide variety of existing data models (as we show in Chapter 3) without requiring complex data structures and nested values, which require a more complex representation and query language.

## 2.3 The ULD Representation Language

In this section, we describe an implementation-independent language for representing configurations in the ULD. By implementation-independent, we mean a language that can be used regardless of the underlying storage of a configuration, for example, a relational database or a set of facts in a deductive database. Thus, the ULD representation language

```

E ::= C | D

C ::= construct c = ( CS | CC | CA | CU )
CS ::= {a1->c1, ..., an->cn} [ F ]
CC ::= ( set | bag | list ) of c1 [ F ]
CA ::= atomic [ F ]
CU ::= c1 | c2 | ... | cn
F ::= conf(domain=x, range=y):c'

D ::= data d = ( DS | DC | DA ) [ I ]
DS ::= c {a1:d1, ..., an:dn}
DC ::= c [d1, ..., dn]
DA ::= c
I ::= d-inst:d1, d2, ..., dn

```

Figure 2.5: The syntax for ULD construct and instance expressions.

can be viewed as a high-level syntax for defining ULD configurations. To illustrate, the following examples (Examples 1, 2, and 3 below) use the ULD representation language to describe simplified versions of the relational, XML, and RDF data model constructs, respectively.<sup>2</sup> (Chapter 3 defines these and other data models in more detail.) Note that there are potentially many ways to describe a data model in the ULD, and these examples show only one choice of representation.

As illustrated in Examples 1, 2, and 3, a construct definition can take one of the following forms. Expressions surrounded by `[]` square brackets are optional and expressions of the form `(a|b|c)` denote a choice between expressions *a*, *b*, or *c*. The syntax for expressing constructs and data instances in the ULD, each represented as a ULD expression *E*, is given in Figure 2.5. Note that the symbols `[]` and `|` are tokens in the language, whereas `[ ]` and `|` are not.

- `construct c = {a1->c1, a2->c2, ..., an->cn} [ conf(domain=x, range=y):c' ]`

This expression defines a tuple-construct *c* as a *ct-inst* of construct type *struct-ct*, where *a<sub>1</sub>* to *a<sub>n</sub>* are distinct strings, *n* ≥ 1, and *c<sub>1</sub>* to *c<sub>n</sub>* and *c'* are construct identifiers.

Each expression *a<sub>i</sub>->c<sub>i</sub>* is called a *component* of the construct *c* where *a<sub>i</sub>* is called the

---

<sup>2</sup>We use `uldValue` and `uldValuetype` as special constructs to denote scalar values and value types, which we describe later.

component *selector*. If the *conf* expression is present, construct instances of  $c$  may conform (that is, be connected by a *d-inst* relationship) to construct instances of  $c'$  according to the domain and range constraints on the expression. If the conformance expression is not present, conformance is not permitted for the construct, that is, there cannot be a *d-inst* relationship for the construct's instances. The cardinality constraints on conformance, shown as  $x$  and  $y$  above, restrict the participation of associated instances in *d-inst* relationships, for both the domain and range of the relationship, to either: exactly one, denoted as 1; zero or one, denoted as ?; zero or more, denoted as \*; or one or more, denoted as +.

- **construct**  $c = ( \text{ set } \mid \text{ bag } \mid \text{ list } ) \text{ of } c_1 [ \text{ conf}(\text{domain}=x, \text{range}=y) : c' ]$   
 This expression defines a collection-construct  $c$  as a *ct-inst* of either *set-ct*, *bag-ct*, or *list-ct*. The definition restricts  $c$ 's construct instances to collections whose members must be instances of the construct  $c_1$ . In addition, for instances of a set construct  $c$ , each member must have a distinct identifier.<sup>3</sup> This construct can also contain a conformance definition.
- **construct**  $c = \text{atomic} [ \text{ conf}(\text{domain}=x, \text{range}=y) : c' ]$   
 This expression defines an atomic type  $c$  as a *ct-inst* of construct type *atomic-ct*. This construct can also contain a conformance definition.
- **construct**  $c = c_1 \mid c_2 \mid \dots \mid c_n$   
 This expression defines a union-construct  $c$  as a *ct-inst* of construct type *union-ct*, where  $c_1$  to  $c_n$  are distinct construct identifiers for  $n \geq 2$  such that, for all  $c_i$ , for  $1 \leq i \leq n$ ,  $c \neq c_i$ . We also require acyclic union definitions, that is, no two constructs can be directly (or indirectly) unions of each other. Each instance of  $c_1$  to  $c_n$  is considered an instance of  $c$ , however, we do not allow instances of only  $c$  directly. A union construct provides a simple mechanism to group heterogeneous structures (for example, atomics and structs) into a single, union type, as opposed

---

<sup>3</sup>We could similarly define a construct type, for example, *val-set*, that would require unique values, using deep equality [48], as opposed to identifiers, using shallow equality [48]. For simplicity, we only give identifier-based sets, but provide a language for expressing other constraints such as value restrictions on sets.

to *isa* relationships, which offer inheritance semantics, and only group like structures (for example, classes).

For the relational data model shown in Example 1, tables and relation types are one-to-one such that each table conforms to exactly one relation type and vice versa. Similarly, each tuple in a table must conform, as shown by the range restriction of 1, to the table's associated relation type. We assume each relation type can have at most one primary key, for simplicity. Note that in the example below, we do not define the constraints implied by a key. In Chapter 3, we discuss the use of data-model constraints in more detail. (Example schema and instance data for the relational data model are presented later, in Example 4.)

**Example 1** (*The relational data model*)

```
% schema constructs
construct relation    = {hasName->uldString, hasAtts->attList}
construct attList     = list of attribute
construct attribute   = {hasName->uldString, hasDomain->uldValuetype}
construct pKey        = {forRel->relation, keyAtts->pKeyAttList}
construct fKey        = {forRel->relation, toRel->relation, keyAtts->fKeyAttList}
construct pKeyAttList = list of attribute
construct fKeyAttList = list of attribute
% data constructs
construct table       = bag of tuple conf(domain=1,range=1):relation
construct tuple       = list of uldValue conf(domain=*,range=1):relation
```

The XML data model [22] shown in Example 2 includes constructs for element types, attribute types, elements, attributes, content models, and content, where element types contain attribute types and content specifications, elements can optionally conform to element types, and attributes can optionally conform to attribute types. We simplify content models to sets of element types for which a conforming element must have at least one subelement for each corresponding type.

**Example 2** (*The XML with DTD data model*)

```
% schema constructs
construct pcdData      = atomic
construct cdata        = atomic
construct elemType     = {hasName->uldString, hasAtts->attDefList,
                          hasModel->contentDef}
construct attDefList = set of attDef
construct attDef      = {hasName->uldString}
construct contentDef = set of elemType
% data constructs
construct element      = {hasTag->uldString, hasAtts->attSet, hasChildren->content}
                          conf(domain=*,range=?):elemType
construct attSet       = set of attribute
construct attribute    = {hasName->uldString, hasVal->cdata}
                          conf(domain=*,range=?):attDef
construct content      = list of node
construct node         = element | pcdData
```

Finally, Example 3 shows the RDF data model with RDF Schema (RDFS) [23, 47] expressed in the ULD. The description includes constructs for classes, properties, resources, and triples. A triple in RDF contains a subject, predicate, and object, where a predicate can be an arbitrary resource, including a defined property. In RDFS, `rdf:type`, `rdfs:subClassOf`, and `rdfs:subPropertyOf` are considered special RDF properties for denoting instance and specialization relationships. However, we model these properties using conformance (for `rdf:type`) and the constructs `subClass` and `subProp`. For example, a subclass relationship is represented by instantiating (using *c-inst*) a `subClass` construct as opposed to using the special `rdfs:subClassOf` RDF property. This approach allows RDF properties and structural relationships to be decoupled. That is, in the ULD representation, RDF properties only model standard relationships instead of being overloaded to represent both standard relationships and type, subclass, and subproperty information. This approach does not limit the expressibility of RDF: partial, optional, and multiple levels of schema are still possible.

**Example 3** (*The RDF(S) data model*)

```

% schema constructs
construct resource = rdfType | simpleRes
construct rdfType = class | prop
construct simpleRes = {hasURI->uldURI} conf(domain=*,range=*) : class
construct class     = {hasURI->uldURI, hasLabel->uldString}
                    conf(domain=*,range=*) : class
construct prop      = {hasURI->uldURI, hasLabel->uldString, hasDomain->class,
                    hasRange->rangeVal} conf(domain=*,range=*) : rdfType
construct rangeVal  = class | uldValueType
construct subClass  = {hasSub->class, hasSuper->class}
construct subProp   = {hasSub->prop, hasSuper->prop}

% data constructs
construct triple    = {hasPred->resource, hasSubj->resource, hasObj->objVal}
construct objVal    = resource | literal
construct literal   = atomic

```

Every ULD configuration can be initialized with default constructs that represent typical primitive value types, such as *String*, *Boolean*, *Integer*, *URI*, etc., as instances of *atomic-ct*. In addition, *uldValueType* and *uldValue*, as shown in Figure 2.6, are special constructs that work together to provide a mechanism for data models to permit user-defined primitive types (for example, to support relational domains). (Note that in Figure 2.6, we shorten *uldValue* to *value* and *uldValueType* to *valuetype*.) The *uldValue* construct is defined as the union (that is, a union construct) of all defined *atomic-ct* constructs. Thus, when a new *atomic-ct* construct is created, it is added to *uldValue*'s definition.<sup>4</sup> The *uldValueType* construct has an instance with the same name, represented as a string value, as each construct of *uldValue*. For example, to add a new primitive type for date values, we would create a new construct with the identifier *date*, assign the construct to *uldValue* (recall *uldValue* is a union construct), and create a new *uldValueType* instance '*date*' (see Figure 2.6).

Sample schema and data (that is, the bottom level of Figure 2.2) for the relational, XML, and RDF data models are shown in Examples 4, 5, and 6, respectively. Example 4

---

<sup>4</sup>We assume that the construct is added by the system managing the ULD configuration, and not by the user.




Figure 2.6: The default `uldValue` and `uldValuetype` constructs for representing scalar types.

gives a subset of the schema and data from Figure 1.1, Example 5 gives a subset of the schema and data from Figure 1.2, and Example 6 gives a subset of the schema and data from Figure 1.3. As shown, expressions for defining construct instances take the form:  $i = c \ v \ \mathbf{d-inst}:i_1, i_2, \dots, i_l$ , where  $c$  is a construct,  $v$  is a valid value (see below) for construct  $c$ , and  $i_1$  to  $i_l$  are construct-instance identifiers for  $l \geq 0$ . The expression defines  $i$  as a *c-inst* of  $c$  with value  $v$ . If  $c$  participates in a union construct  $c'$ ,  $i$  is considered a *c-inst* of both  $c$  and  $c'$ . Further,  $i$  is a *d-inst* of each construct instance in  $i_1$  to  $i_l$ , each of which must be a *c-inst* of the construct that  $c$  conforms to. The **d-inst** keyword is omitted if  $i$  does not participate in a *d-inst* relationship with another construct instance (that is,  $l = 0$ ).

As described above, a syntactically well-formed construct instance is allowed in a configuration only if it satisfies the constraints of its associated construct. We end this section by giving an informal definition of when a value is valid (with respect to its associated construct).

Given the construct  $c$ , we write  $c \in \text{ct-extent-of}(ct)$  if and only if  $c$  is defined as a *ct-inst* of construct type  $ct$ . Let  $\text{ct-extent-of}(ct)$  be the set of all instances in a configuration of construct type  $ct$ . Similarly, let  $c\text{-extent-of}(c)$  be the set of all construct instances in a configuration that are a *c-inst* of construct  $c$ , and  $d\text{-extent-of}(d)$  be the set of all construct instances in a configuration that are a *d-inst* of  $d$ . The expressions  $i \in c\text{-extent-of}(c)$  and  $i \in d\text{-extent-of}(i')$  are true if and only if  $i$  is defined as a *c-inst* of construct  $c$  and a

$d-inst$  of  $i'$ . The relation  $conf(c, c', x, y)$  is true if and only if a conformance relationship is defined from  $c$  to  $c'$  (that is, instances of construct  $c$  can be a  $d-inst$  of  $c'$ ) with a domain constraint of  $x$  and a range constraint of  $y$ . The following constraints are imposed by constructs on their instances and must hold for a value to be valid. We note that each construct instance can only be defined once in a configuration, and must have exactly one associated, non-union construct, that is, it must be a  $c-inst$  of exactly one construct that is not a  $ct-inst$  of  $union-ct$ .

- **Structural Constraints.** Let  $i$  be a construct instance such that  $i \in c-extent-of(c)$ .

If  $c$  is a tuple construct (that is,  $c \in ct-extent-of(struct-ct)$ ), the value of  $i$  must contain the same number of components as  $c$ , each component selector of  $i$  must be a distinct selector of  $c$ , and each component value of  $i$  must be an instance of the associated component construct value for  $c$ . (Note that we require all components to have a non-null value.) If  $c$  is a collection construct (that is,  $c \in ct-extent-of(set-ct) \cup ct-extent-of(bag-ct) \cup ct-extent-of(list-ct)$ ), the value of  $i$  must be a collection whose members are instances of the construct that  $c$  is a collection of, and if  $c$  is a set,  $i$ 's elements must have distinct identifiers. If  $c$  is an atomic construct (that is,  $c \in ct-extent-of(atomic-ct)$ ), we assume that the value of  $i$  satisfies  $c$ . For example, if  $c$  is a string,  $i$  is a well-formed string; if  $c$  is an integer,  $i$  is a well-formed integer; and so on. We note that, if  $c$  is a union construct (that is,  $c \in ct-extent-of(union-ct)$ ),  $c$  itself does not impose any additional restrictions on  $i$ 's value.

- **Conformance Constraints.** If  $i$  is an instance of construct  $c$  (that is,  $i \in c-extent-of(c)$ ),  $i'$  is an instance of construct  $c'$  (that is,  $i' \in c-extent-of(c')$ ), and  $c$  can conform to  $c'$  (that is,  $conf(c, c', x, y)$  is true for some  $x, y$ ),  $i$  is allowed to be a  $d-inst$  of  $i'$  (that is, we allow  $i \in d-extent-of(i')$ ). Further, domain and range cardinality constraints on conformance restrict the allowable  $d-inst$  relationships as follows. If the domain constraint on the conformance relationship is  $+$  (that is,  $x = +$ ), each  $i'$  in the configuration must participate in at least one  $d-inst$  relationship, where there is at least one  $i$  for each  $i'$  such that  $i \in d-extent-of(i')$ . If the domain constraint on the conformance relationship is  $1$ , each  $i'$  in the configuration must have exactly one

$i$  such that  $i \in d\text{-extent-of}(i')$ . And, if the domain constraint on the conformance relationship is  $?$ , each  $i'$  in the configuration has at most one  $i$  such that  $i \in d\text{-extent-of}(i')$ . Similarly, if the range constraint on the conformance relationship is  $+$  (that is,  $y = +$ ), every  $i$  in the configuration must be the  $d\text{-inst}$  of at least one  $i'$ , that is, there must be at least one  $i'$  such that  $i \in d\text{-extent-of}(i')$ . If the range constraint on the conformance relationship is  $1$ , each  $i$  in the configuration must be a  $d\text{-inst}$  of exactly one  $i'$ . Finally, if the range constraint on the conformance relationship is  $?$ , each  $i$  can be the  $d\text{-inst}$  of at most one  $i'$ .

**Example 4** (*Sample relational data and schema*)

```
% schema
movie = relation {hasName:'movie', hasAtts:a11} a11 = attList
[a1,a2,a3,a4,a5] a1 = attribute {hasName:'mid', domain:'uldInteger'}
a2 = attribute {hasName:'title', domain:'uldString'} a3 = attribute
{hasName:'runtime', domain:'uldInteger'} a4 = attribute
{hasName:'genre', domain:'uldString'} a5 = attribute
{hasName:'company', domain:'uldString'} key1 = pKey {forRel:movie,
keyAtts:pkal1} pkal1 = pKeyAttList [a1] cast = relation
{hasName:'cast', atts:a13} a13 = attlist [a6,a7,a8] a6 = attribute
{hasName:'mid', domain:'uldInteger'} a7 = attribute {hasName:'actor',
domain:'uldString'} a8 = attribute {hasName:'character',
domain:'uldString'}
% data
tbl1    = table [t1] d-inst:movie
t1      = tuple [1,'The Usual Suspects',106,'thriller','Gramercy']
          d-inst:movie
tbl2    = table [t2] d-inst:cast
t2      = tuple [1,'Roger Verbal Kint','Kevin Spacey'] d-inst:cast
```

**Example 5** (*Sample XML with DTD data and schema*)

```
% schema
movie   = elemType {hasName:'movie', hasAtts:nilad, hasModel:moviecm}
Nilad   = attDefSet []
moviecm = contentDef [title, studio, genre, actor]
```

```

title = elemType {hasName:'title', hasAtts:nilad, hasModel:nilcm}
nilcm = contentDef []
genre = elemType {hasName:'genre', hasAtts:nilad, hasModel:nilcm}
actor = elemType {hasName:'actor', hasAtts:actorat, hasModel:nilcm}
actorat = attDefList [role]
role = attdef {name:'role'}
% data
e1 = element {hasTag:'movie', hasAtts:nilas, hasChildren:e1cnt}
      d-inst:movie
nilas = attSet []
e1cnt = content [e2,e3,e4]
e2 = element {hasTag:'title', hasAtts:nilas, hasChildren:e2cnt}
      d-inst:title
e2cnt = content ['The Usual Suspects']
e3 = element {hasTag:'genre', hasAtts:nilas, hasChildren:e3cnt}
      d-inst:genre
e3cnt = content ['thriller']
e4 = element {hasTag:'actor', hasAtts:e4as, hasChildren=e4cnt}
      d-inst:actor
e4as = attSet [a1]
e4cnt = content ['Kevin Spacey']
a1 = attribute {hasName:'role', hasVal:'supporting'} d-inst:role

```

**Example 6** (*Sample RDF(S) data and schema*)

```

% schema
film = class {hasURI:'#film', hasLabel:'film'}
title = prop {hasURI:'#title', hasLabel:'hasTitle', hasDomain:film,
              hasRange:'literal'}
thriller = class {hasURI:'#thriller', hasLabel:'thriller'}
filmthril = subclass {hasSub:thriller, hasSuper:film}
% data
f1 = simpleRes {hasURI:'http://.../review.html'} d-inst:thriller,film
t1 = triple {hasPred:title, hasSubj:f1, hasObj:'The Usual Suspects'}

```

## 2.4 The ULD Query Language

This section describes a query language for the ULD that exploits the capabilities of Datalog [2] (we describe Datalog in more detail below). A ULD query is expressed as a

Datalog program, that is, a set of restricted horn-clause rules, and is executed against a particular configuration. As an example, the following query finds all available class names within an RDF configuration. Note that upper-case terms denote variables and lower-case terms denote constants. Also note that the expression `C.hasLabel:X`, which is non-standard Datalog syntax, selects the `hasLabel` component-value `X` from the class-construct instance `C`. We describe additional extensions to Datalog in Section 4.2.2.

$$\text{classname}(X) \leftarrow \text{c-inst}(C, \text{class}), C.\text{hasLabel}:X.$$

The rule is read as “If `C` is an RDF class and the label of `C` is `X`, then `X` is a classname.” As another example, the following query returns the property names of all classes in an RDF configuration. Note that this query, like the previous one, is expressed solely against the schema of the source, that is, we are only accessing classes and properties.

$$\begin{aligned} \text{hasProp}(X, Y) \leftarrow & \text{c-inst}(C, \text{class}), \text{c-inst}(P, \text{prop}), P.\text{hasDomain}:C, C.\text{hasLabel}:X, \\ & P.\text{hasLabel}:Y. \end{aligned}$$

After finding the available RDF classes and their corresponding properties, we can then use this information to find corresponding data instances. That is, a user could issue the previous query, see that the source contains title properties, and then construct the following query to return all title values in the configuration.

$$\begin{aligned} \text{filmtitle}(X) \leftarrow & \text{c-inst}(P, \text{prop}), P.\text{hasLabel}:\text{'title'}, \text{c-inst}(T, \text{triple}), T.\text{hasPred}:P, \\ & T.\text{hasObj}:X. \end{aligned}$$

With the ULD, it is possible to query data (including accessing data without first accessing schema), schema (as shown in the previous queries), and data-model constructs. For example, the following query is expressed directly against data, and returns the URI of all RDF resources used as a property in at least one triple (note that the resource may or may not be associated with schema).

$$\text{dataprop}(X) \leftarrow \text{c-inst}(T, \text{triple}), T.\text{hasPred}:P, P.\text{hasURI}:X.$$

Once this query is executed, we may wish to find additional information about a particular resource. For example, the following query returns all values of the title resource used as the property of a triple. (Note that, in RDF, resources are typically represented with a leading '#' symbol.)

$$\text{propval}(X) \leftarrow \text{c-inst}(T, \text{triple}), T.\text{hasPred}:P, P.\text{hasURI}:\text{'\#title'}, T.\text{hasObj}:X.$$

The following queries are similar to the previous examples, but are expressed against an XML configuration. The first query finds the names of all available element types in the source, the second finds, for each element-type name, its corresponding attribute-definition names, the third finds the set of movie titles, the fourth finds all available attribute names as a data query, and the last query finds the set of attribute values for title attributes in the configuration. The membership operator  $\in$  in the second query is used to access elements within a collection.

$$\begin{aligned} \text{elemtype}(X) &\leftarrow \text{c-inst}(E, \text{elemType}), E.\text{hasName}:X. \\ \text{atttype}(X,Y) &\leftarrow \text{c-inst}(E, \text{elemType}), E.\text{hasName}:X, E.\text{hasAtts}:AL, AT \in AL, \\ &\quad AT.\text{hasName}:Y. \\ \text{title}(X) &\leftarrow \text{c-inst}(AT, \text{attDef}), AT.\text{hasName}:\text{'title'}, d\text{-inst}(A, AT), \\ &\quad A.\text{hasVal}:X. \\ \text{atts}(X) &\leftarrow \text{c-inst}(A, \text{attribute}), A.\text{hasName}:X. \\ \text{titleattval}(X) &\leftarrow \text{c-inst}(A, \text{attribute}), A.\text{hasName}:\text{'title'}, A.\text{hasVal}:X. \end{aligned}$$

Finally, queries can be posed against data-model constructs directly, for example, to find the constructs used by an information source. The following query returns all constructs that serve as *struct-ct* schema constructs and their component selectors. Note that the expression  $SC.P \rightarrow C$  is similar to the expression  $X.P:Y$ , but applies to constructs instead of construct instances.

$$\text{schemastruct}(SC, P) \leftarrow \text{ct-inst}(SC, \text{struct-ct}), \text{conf}(DC, SC, X, Y), SC.P \rightarrow C.$$

Similarly, it is possible to access data values within a configuration, without knowledge of the data model. For example, the following query returns all primitive values of *struct-ct* constructs that serve as data for a corresponding schema item within a configuration.

$$\begin{aligned} \text{dataval}(V) \leftarrow & \text{ct-inst}(\text{DC}, \text{struct-ct}), \text{conf}(\text{DC}, \text{SC}, X, Y), \text{DC.S} \rightarrow \text{C}, \\ & \text{ct-inst}(\text{C}, \text{atomic-ct}), \text{c-inst}(\text{D}, \text{DC}), \text{D.S:V}. \end{aligned}$$

The remainder of this section describes, in more detail, the use of Datalog for expressing queries over ULD configurations. We begin by giving a brief introduction to Datalog in Section 2.4.1, then describe an approach for interpreting a ULD configuration as a source for Datalog queries in Section 2.4.2. In Section 2.5 we give a method for defining configuration constraints, which are expressed as Datalog programs.

### 2.4.1 Datalog

Datalog is a logic-based language for querying relational databases. A Datalog query (or program) consists of a finite set of *Datalog rules*. Each Datalog rule takes the following form, for  $n \geq 1$ .

$$R_1(u_1) \leftarrow R_2(u_2), R_3(u_3), \dots, R_n(u_n).$$

This rule is read as “ $R_2(u_2)$  and  $R_3(u_3)$  and ... and  $R_n(u_n)$  implies  $R_1(u_1)$ .” In a Datalog rule, each  $R_i$  represents a relation and each  $u_i$  represents a term of  $R_i$ . Often,  $R_i$  is referred to as a *predicate* applied to terms  $u_i$ , where the expression  $R_i(u_i)$  is referred to as an *atomic formula* (or just *formula*) and  $u_i$  consists of a list of *terms*  $u_i = u_{i1}, u_{i2}, \dots, u_{im}$  for  $m \geq 1$ . We use the standard convention where predicates and constants begin with lower-case letters (or possibly special characters) and variables begin with upper-case letters. Datalog rules are *range restricted*, that is, each variable appearing in  $u_1$  must appear in  $u_2, \dots, u_n$ . The *head* of the rule occurs on the left side of the arrow and always contains a single formula  $R_1(u_1)$ . The *body* of the rule consists of the set of conjoined formulas  $R_2(u_2), \dots, R_n(u_n)$ .

A Datalog program contains finite sets of *extensional* and *intensional* relations, whose instances (that is, sets of associated facts) are called the *extensional database* (EDB) and *intensional database* (IDB), respectively. An extensional relation is a relation that exists in the underlying database and can only occur in the body of a rule. Thus, an extensional formula  $R_i(u_i)$  corresponds to the relation in the underlying database with the name  $R_i$ ,

and  $u_i$  represents a tuple (as a list of terms) with the same arity as the corresponding database relation. An intensional relation is a relation occurring in the head of some rule. The result of a executing a Datalog program is an IDB.

Equality and inequality predicates ( $=$ ,  $\neq$ ,  $<$ ,  $>$ , and so on) are treated as special relations in Datalog. However, the introduction of equality further constrains the range-restriction requirement. Namely, in every equality expression containing two variables, at least one variable must occur as a variable in a relational atom  $R_i(u_i)$  in the body of the rule. For an inequality, both variables must occur in a relational atom in the body of the rule. We note that equality and inequality relations are typically represented using infix notation, that is,  $x = y$  instead of  $=(x, y)$ . Similar to equality and inequality relations, Datalog can also be extended to support functions [80]. The ULD query language uses functions and relations to access configurations, which we describe further in Section 2.4.2.

In the *minimal model semantics*, the answer to a Datalog program  $P$  over input  $I$  (where  $I$  represents the given instances, or tuples, of the extensional relations) is defined as the minimal model of  $P$  containing  $I$ . Alternatively, the *least fixpoint semantics* states that an answer to  $P$  is the least fixpoint of  $P$ , which is computed by repeatedly applying the rules of  $P$  on input  $I$ . Every Datalog program, as defined above, has a least fixpoint, which is also the minimal model.

As defined, only positive (that is, non-negated) formulas are permitted in a Datalog rule. Datalog $\neg$  extends Datalog by allowing negative formulas in the body of rules.<sup>5</sup> Similar to Datalog, every variable in the head of a *range-restricted* Datalog $\neg$  rule appears in some positive relation occurrence in the body. In addition, every variable in a negated formula of a Datalog $\neg$  rule must occur in a positive formula in the body of the rule. We note that not all Datalog $\neg$  programs have a unique least fixpoint.

A restricted form of Datalog $\neg$  considers only *stratified* programs, which place a syntactic restriction on rules. (The ULD query language assumes queries are expressed as stratified Datalog $\neg$  programs.) Namely, a stratified program can be partitioned into a finite number of sets (or layers) of rules  $P^1, P^2, \dots, P^n$  such that (1) for each IDB relation

---

<sup>5</sup>Note that Datalog $\neg$  uses the closed-world assumption, where facts that are not explicitly stated as true or that cannot be inferred are assumed to be false.

$R$ , all rules in  $P$  defining  $R$  are in the same partition, (2) if  $R(u) \leftarrow \dots R'(v) \dots$  is a rule in  $P$ ,  $R'$  is an IDB relation,  $R'$  is in partition  $P^i$ , and  $R$  is in partition  $P^j$ , then  $i \leq j$ , and (3) if  $R(u) \leftarrow \dots \neg R'(v) \dots$  is a rule in  $P$ ,  $R'$  is an IDB relation,  $R'$  is in partition  $P^i$ , and  $R$  is in partition  $P^j$ , then  $i < j$ . Stratification gives a natural order for executing rules, namely, the fixpoint of partition  $P^1$  is computed first, followed by the fixpoint of partition  $P^2$ , and so on, where partition  $P^n$  is computed last. A program that can be expressed using stratified Datalog $\neg$  is guaranteed to have a unique least fixpoint. However, not all Datalog $\neg$  programs can be stratified.

Finally, we note that it is possible to construct an unsatisfiable query in Datalog $\neg$ , that is, regardless of the EDB, the query always has an empty result. For example, the rule  $p(X) \leftarrow q(X), \neg q(X)$  will never return an  $X$  that makes  $p(X)$  true.

#### 2.4.2 Querying ULD Configurations using Datalog

As mentioned in the beginning of this section, a ULD query is expressed using Datalog, and more precisely, using stratified Datalog $\neg$ . Thus, we treat a configuration  $\mathcal{F}$  as a database instance consisting of a fixed set of extensional relations. To describe a configuration as an EDB, we represent a configuration  $\mathcal{F}$  as a tuple  $(\mathcal{I}_{CT}, \mathcal{I}_C, \mathcal{I}_D, \mathcal{V}, val, ct\text{-extent-of}, c\text{-extent-of}, d\text{-extent-of}, conf)$  as follows.

- $\mathcal{I}_{CT}$ ,  $\mathcal{I}_C$ , and  $\mathcal{I}_D$  are sets that consist, respectively, of construct-type identifiers, construct identifiers, and construct-instance identifiers. We require  $\mathcal{I}_{CT}$ ,  $\mathcal{I}_C$ , and  $\mathcal{I}_D$  to be pairwise disjoint sets.
- $\mathcal{V}$  is the set of values of construct and data identifiers.
- $val : \mathcal{I}_C \cup \mathcal{I}_D \rightarrow \mathcal{V}$  maps each construct and data identifier to a value.
- $ct\text{-extent-of} : \mathcal{I}_{CT} \rightarrow \mathcal{P}(\mathcal{I}_C)$  maps construct-type identifiers to their associated construct-instance identifiers, where  $\mathcal{P}(\mathcal{I}_C)$  denotes the powerset (set of subsets) of  $\mathcal{I}_C$ .
- $c\text{-extent-of} : \mathcal{I}_C \rightarrow \mathcal{P}(\mathcal{I}_D)$  maps construct identifiers to their associated construct-instance identifiers.

- *d-extent-of*:  $\mathcal{I}_{\mathcal{D}} \rightarrow \mathcal{P}(\mathcal{I}_{\mathcal{D}})$  maps construct-instance identifiers to their associated construct-instance identifiers.
- *conf*:  $\mathcal{I}_{\mathcal{C}} \times \mathcal{I}_{\mathcal{C}} \times \{*, +, ?, 1\} \times \{*, +, ?, 1\} \rightarrow \text{Boolean}$  is the conformance relation, such that given an input  $\langle c, c', x, y \rangle$  for  $c, c' \in \mathcal{I}_{\mathcal{C}}$ , *conf* returns true if  $c$  can conform to  $c'$  with domain constraint  $x$  and range constraint  $y$ .

The values  $\mathcal{V}$  of  $\mathcal{F}$  are represented as follows.

- The tuple  $\langle (s_1, c_1), (s_2, c_2), \dots, (s_n, c_n) \rangle$  is in  $\mathcal{V}$  for every construct identifier  $c \in \text{ct-extent-of}(\text{struct-ct})$  such that  $n$  is the number of components of  $c$ ,  $s_i$  is the selector of the  $i$ -th component of  $c$ , and  $c_i$  is the construct identifier of the  $i$ -th component of  $c$ , for  $1 \leq i \leq n$ .
- The tuple  $\langle (s_1, v_1), (s_2, v_2), \dots, (s_n, v_n) \rangle$  is in  $\mathcal{V}$  for every construct-instance identifier  $d \in \mathcal{I}_{\mathcal{D}}$  such that there is a  $c \in \text{ct-extent-of}(\text{struct-ct})$  where  $d \in \text{c-extent-of}(c)$ , and  $n$  is the number of components of  $d$ ,  $s_i$  is the selector of the  $i$ -th component of  $d$ , and  $v_i$  is the value of the  $i$ -th component of  $d$ , for  $1 \leq i \leq n$ .
- The singleton  $c'$  is in  $\mathcal{V}$  for every construct identifier  $c \in \text{ct-extent-of}(\text{set-ct}) \cup \text{ct-extent-of}(\text{list-ct}) \cup \text{ct-extent-of}(\text{bag-ct})$  such that  $c$  is defined as a collection of  $c'$ .
- The set  $\{c_1, c_2, \dots, c_n\}$  is in  $\mathcal{V}$  for every union construct identifier  $u \in \text{ct-extent-of}(\text{union-ct})$  such that  $u$  is defined as a union of construct identifiers  $c_1, c_2, \dots, c_n$ .
- The set  $\{v_1, v_2, \dots, v_n\}$  is in  $\mathcal{V}$  for every construct-instance identifier  $v \in \text{c-extent-of}(c)$  such that  $c \in \text{ct-extent-of}(\text{set-ct})$  and  $v$  is defined as the set of construct-instance identifiers  $v_1, v_2, \dots, v_n$ .
- The bag  $[v_1, v_2, \dots, v_n]$  is in  $\mathcal{V}$  for every construct-instance identifier  $v \in \text{c-extent-of}(c)$  such that  $c \in \text{ct-extent-of}(\text{bag-ct})$  and  $v$  is defined as the bag of construct-instance identifiers  $v_1, v_2, \dots, v_n$ .

- The tuple  $\langle v_1, v_2, \dots, v_n \rangle$  is in  $\mathcal{V}$  for every construct-instance identifier  $v \in c\text{-extent-of}(c)$  such that  $c \in ct\text{-extent-of}(list\text{-}ct)$  and  $v$  is defined as a list of construct-instance identifiers where  $v_1$  is at index 1,  $v_2$  is at index 2, ..., and  $v_n$  is at index  $n$ .
- The singleton  $v$  is in  $\mathcal{V}$  for every construct-instance identifier  $v \in c\text{-extent-of}(c)$  such that  $c \in ct\text{-extent-of}(atomic\text{-}ct)$ . Note that we assume an atomic value and its associated identifier are equivalent (in this case,  $v$ ).

We assume that standard functions and relations are defined for the types of  $\mathcal{V}$ . That is,  $v.s$  returns the corresponding construct or construct-instance identifier given a tuple  $v$  and a selector  $s$  (where  $s$  is a selector for  $v$ ). The membership relation  $\in$  is defined for sets, lists, and bags. The length-function  $|v|$  takes a collection  $v$  and returns the number of elements of  $v$ . The function  $v[i]$  returns the item at the  $i$ th position of  $v$ , where  $v$  is a list and  $i$  is a positive integer between 1 and  $|v|$ . And the function  $x \in\in v$  returns the cardinality of  $x$  in  $v$ .

Finally, for a configuration  $\mathcal{F}$ , the interpretation  $I_{\mathcal{F}}$  consists of a non-empty set  $\Delta^I$  (the domain of the interpretation) and maps formulas in the query language to expressions in the configuration  $\mathcal{F}$ . We call  $I_{\mathcal{F}}$  the *database instance mapping* of configuration  $\mathcal{F}$ . Note that we use the standard Datalog semantics, thus, we only describe the interpretation of atomic formulas for the purpose of determining truth assignments, and not arbitrary Datalog expressions. For example, the standard interpretation (for Datalog) of a relational database maps each constant to a constant in the domain of the database, and each formula  $R(x_1, \dots, x_n)$  for constants  $x_1, \dots, x_n$  to the corresponding tuple  $\langle x_1, \dots, x_n \rangle$  of the relation instance  $R$  [73]. Given an interpretation  $I_{\mathcal{F}}$  for the configuration  $\mathcal{F}$ , we define the domain of the interpretation  $\Delta^I$  as follows. (Note that atomic values are given in  $\mathcal{I}_{\mathcal{D}}$ . Also, the second line adds component-selector definitions used in constructs, whereas the third line adds component selectors used in construct instances.)

$$\begin{aligned}
\Delta^I = & \mathcal{I}_{\mathcal{CT}} \cup \mathcal{I}_{\mathcal{C}} \cup \mathcal{I}_{\mathcal{D}} \cup \\
& \{s \mid \exists c (c \in ct\text{-extent-of}(struct\text{-}ct) \wedge val(c) = \langle \dots, (s, c'), \dots \rangle)\} \cup \\
& \{s \mid \exists d \exists c (d \in c\text{-extent-of}(c) \wedge val(d) = \langle \dots, (s, c'), \dots \rangle)\} \cup \\
& \{n \mid \exists d \exists c (d \in c\text{-extent-of}(c) \wedge c \in ct\text{-extent-of}(set\text{-}ct) \wedge |val(d)| = n)\} \cup \\
& \{n \mid \exists d \exists c (d \in c\text{-extent-of}(c) \wedge c \in ct\text{-extent-of}(list\text{-}ct) \wedge |val(d)| = n)\} \cup
\end{aligned}$$

$$\begin{aligned}
& \{n \mid \exists d \exists c (d \in c\text{-extent-of}(c) \wedge c \in ct\text{-extent-of}(bag\text{-}ct) \wedge |val(d)| = n)\} \cup \\
& \{i \mid \exists d \exists c (d \in c\text{-extent-of}(c) \wedge c \in ct\text{-extent-of}(list\text{-}ct) \wedge 1 \leq i \leq |val(d)|)\} \cup \\
& \{i \mid \exists d \exists c (d \in c\text{-extent-of}(c) \wedge c \in ct\text{-extent-of}(bag\text{-}ct) \wedge 1 \leq i \leq |val(d)|)\}
\end{aligned}$$

The interpretation  $I_{\mathcal{F}}$  is given in Figure 2.7 (for an arbitrary  $\mathcal{F}$ ). Each expression gives a mapping such that the expression on the left side of the equality sign is true if and only if the expression on the right side of the equality sign is true.

$$\begin{aligned}
ct\text{-inst}(x, y)^I & \equiv x \in ct\text{-extent-of}(y) \\
c\text{-inst}(x, y)^I & \equiv x \in c\text{-extent-of}(y) \\
d\text{-inst}(x, y)^I & \equiv x \in d\text{-extent-of}(y) \\
conf(x, y, d, r)^I & \equiv conf(x, y, d, r) \\
(x.s \rightarrow y)^I & \equiv x \in ct\text{-extent-of}(struct\text{-}ct) \wedge val(x).s = y \\
(x.s.y)^I & \equiv \exists(c) x \in c\text{-extent-of}(c) \wedge c \in ct\text{-extent-of}(struct\text{-}ct) \wedge val(x).s = y \\
(y \in x)^I & \equiv \exists(c) x \in c\text{-extent-of}(c) \wedge [c \in ct\text{-extent-of}(set\text{-}ct) \vee \\
& \quad c \in ct\text{-extent-of}(list\text{-}ct) \vee c \in ct\text{-extent-of}(bag\text{-}ct)] \wedge y \in val(x) \\
(x[i] = y)^I & \equiv \exists(c) x \in c\text{-extent-of}(c) \wedge c \in ct\text{-extent-of}(list\text{-}ct) \wedge val(x)[i] = y \\
(|x| = n)^I & \equiv \exists(c) x \in c\text{-extent-of}(c) \wedge [c \in ct\text{-extent-of}(set\text{-}ct) \vee \\
& \quad c \in ct\text{-extent-of}(list\text{-}ct) \vee c \in ct\text{-extent-of}(bag\text{-}ct)] \wedge |val(x)| = n \\
(y \in x = n)^I & \equiv \exists(c) x \in c\text{-extent-of}(c) \wedge c \in ct\text{-extent-of}(list\text{-}ct) \wedge y \in val(x) = n \\
setof(x, y)^I & \equiv x \in ct\text{-extent-of}(set\text{-}ct) \wedge val(x) = y \\
bagof(x, y)^I & \equiv x \in ct\text{-extent-of}(bag\text{-}ct) \wedge val(x) = y \\
listof(x, y)^I & \equiv x \in ct\text{-extent-of}(list\text{-}ct) \wedge val(x) = y \\
unionof(x, y)^I & \equiv x \in ct\text{-extent-of}(union\text{-}ct) \wedge y \in val(x)
\end{aligned}$$

Figure 2.7: The interpretation of ULD configurations.

## 2.5 The ULD Constraint Language

The ability to query the ULD is important for providing uniform access to information sources. The ULD query language can also be used to specify additional constraints on configurations. In this section, we describe a special syntax for expressing constraints based on the ULD query language. Our goal is to define a language that can be used to describe common ULD constraints, while exploiting the expressive power of Datalog. Therefore, we have chosen a syntax for expressing ULD constraints that can be directly converted to stratified Datalog $\neg$  programs. We also show how such constraints can be used to further refine data models, and in particular, to express conformance constraints beyond simple cardinality restrictions. In general, further refining a data model by providing additional constraints using the language presented here is optional. However, we require

domain and range cardinality restrictions on conformance, for example, so that additional tools can determine whether schema is required or optional.

Constraints against the ULD are represented as stratified Datalog $\neg$  programs with a specific set of formulas. The result of executing the program against a configuration is used to determine whether the constraint is satisfied. To make constraints easier to specify, we introduce the following syntax for defining ULD constraints, where the name of the constraint is represented as  $t$ ,  $n \geq 1$ , and  $m \geq 0$ .

constraint  $t$ :

$$e_1, e_2, \dots, e_n \Leftarrow b_c.$$

$$r_1 \leftarrow b_1.$$

$$r_2 \leftarrow b_2.$$

$$\vdots$$

$$r_m \leftarrow b_m.$$

A constraint expressed with this syntax can be directly converted to a corresponding stratified Datalog $\neg$  program, which we show below. Due to the underlying conversion, we require that  $t$ -body and  $t$ -sat (which are introduced by the conversion) are not used as predicate names in  $r_1$  to  $r_m$  below, where  $t$  is replaced by the name of the constraint. All ULD constraints consist of exactly one rule called the *constraint assertion*, plus an optional set of rules forming a stratified Datalog $\neg$  program, called the *constraint program*. The constraint assertion is shown as the first rule above and the remaining rules define the constraint program. As shown, constraint-assertion rules are distinguished from constraint-program rules by the  $\Leftarrow$  implication symbol. The head of a constraint assertion has one or more extensional relations  $e_1$  to  $e_n$ , and a standard stratified Datalog $\neg$  rule body  $b_c$  containing formulas with either extensional relations or intensional relations taken from  $r_1$  to  $r_m$ . Unlike Datalog, the head of a constraint assertion can contain multiple formulas, positive formulas with variables not mentioned in the body of the rule, and negative formulas. We require all variables occurring in a negative formula in the head of an assertion to occur in a positive formula in the assertion (either the head or the body).

A constraint assertion serves to define the purpose of the constraint, and the constraint

program serves to define temporary, intensional relations used by the constraint assertion. The meaning of a constraint assertion is defined as follows. Let  $v$  be a function that maps a set of variable names to variable assignments (that is, values) and let  $var$  be a function that takes one or more formulas and returns the corresponding set of variables. The constraint program is executed first. The body  $b_c$  of the constraint assertion is then evaluated against both the configuration facts and the resulting IDB, generated from executing the constraint program. Each evaluation of  $b_c$  results in a variable assignment  $v_b$  that makes  $b_c$  true with respect to the configuration and the IDB. Given a formula  $e$ , we write  $e/v$  to denote a new formula with every variable in  $e$  replaced by its corresponding value in the variable assignment  $v$ . Similarly, for variable assignments  $v_1$  and  $v_2$ ,  $e/v_1/v_2$  denotes the formula with every variable in  $e/v_1$  replaced by its corresponding value in the variable assignment  $v_2$ , and so on. Given a configuration  $\mathcal{F}$  with corresponding database instance mapping  $I_{\mathcal{F}}$  and constraint  $t$ , we use the standard notation  $I_{\mathcal{F}} \models t$  whenever  $I_{\mathcal{F}}$  satisfies constraint  $t$ . For a configuration instance  $I_{\mathcal{F}}$  and constraint  $t$ ,  $I_{\mathcal{F}} \models t$  if and only if, for each variable assignment  $v_b$  of  $b_c$ , the following conditions are true.

1. If  $e_i$  in  $e_1, e_2, \dots, e_n$  is a positive formula with every variable in  $var(e_i)$  occurring in  $var(b_c)$  (that is,  $var(e_i) \subseteq var(b_c)$ ), then  $e_i/v_b \in I_{\mathcal{F}}$ .
2. If  $e_i$  in  $e_1, e_2, \dots, e_n$  is a negative formula of the form  $\neg f_i$  with every variable in  $var(f_i)$  occurring in  $var(b_c)$  (that is,  $var(f_i) \subseteq var(b_c)$ ), then  $f_i/v_b \notin I_{\mathcal{F}}$ .
3. If  $E$  is the largest subset of formulas  $e_1$  to  $e_n$  such that every  $e_j \in E$  has at least one variable in  $var(e_j)$  not occurring in  $var(b_c)$  (that is,  $\exists V (V \in var(e_j) \wedge V \notin var(b_c))$ ), then there must exist at least one variable assignment  $v_e$  for  $E$  that makes each formula in  $E$  true for the configuration. That is, for  $E = \{e_{j1}, e_{j2}, \dots, e_{jk}\}$  and  $var(E) - var(b_c) = \{V_1, V_2, \dots, V_l\}$ , the set  $E$  represents the first-order logic expression  $\exists V_1 \exists V_2 \dots \exists V_l (e_{j1}/v_b, e_{j2}/v_b, \dots, e_{jk}/v_b)$ . A variable assignment  $v_e$  makes  $E$  true with respect to the configuration if and only if for every formula  $e_j \in E$ ,  $e_j/v_b/v_e \in I_{\mathcal{F}}$ . (Note that according to the definition of the ULD constraint syntax,  $E$  can only contain positive formulas.)

The ULD constraint language is meant to simplify the task of expressing constraints against ULD configurations. Every well-formed ULD constraint can be rewritten as a stratified Datalog $\neg$  program as follows, where  $t$  is the name of the constraint and both  $t$ -body and  $t$ -sat are intensional relations, introduced as part of the re-writing process.

$$\begin{aligned} t\text{-sat}(var(b_c)) &\leftarrow t\text{-body}(var(b_c)), e_1, e_2, \dots, e_n. \\ t\text{-body}(var(b_c)) &\leftarrow b_c. \\ r_1 &\leftarrow b_1. \\ r_2 &\leftarrow b_2. \\ &\vdots \\ r_m &\leftarrow b_m. \end{aligned}$$

Note that we use the syntax  $p(var(b_c))$  to denote a formula  $p(V_1, V_2, \dots)$  where  $var(b_c) = \{V_1, V_2, \dots\}$ . If  $var(b_c)$  is empty (that is, there are no variables in the body of the constraint assertion), we substitute  $var(b_c)$  with an arbitrary constant symbol (for example, we could choose a constant from one of the formulas of  $b_c$ ). The use of a constant symbol is needed because Datalog requires at least one attribute in every relation.

With this Datalog $\neg$  representation of a ULD constraint, it is straightforward to check whether the constraint is satisfied by a configuration (note that we assume configurations are constrained to be finite). Namely,  $I_{\mathcal{F}} \models t$  if and only if for each fact in the evaluation of the formula  $t$ -sat there is a corresponding fact in the evaluation of  $t$ -body. That is, for each tuple of  $t$ -sat there is an identical tuple in  $t$ -body and vice versa. For example, consider the following constraint (Example 7), which states that, within an RDF configuration, no two classes may have the same label.

**Example 7** (*RDF class-label constraint*)

constraint key1:

$$C1=C2 \Leftarrow \text{c-inst}(C1, \text{class}), \text{c-inst}(C2, \text{class}), C1.\text{hasLabel:L}, C2.\text{hasLabel:L}.$$

The RDF class-label constraint can be directly re-written into the following Datalog rules.

$$\begin{aligned} t\text{-sat}(C1, C2, L) &\leftarrow t\text{-body}(C1, C2, L), C1=C2. \\ t\text{-body}(C1, C2, L) &\leftarrow \text{c-inst}(C1, \text{class}), \text{c-inst}(C2, \text{class}), C1.\text{hasLabel:L}, C2.\text{hasLabel:L}. \end{aligned}$$

Assume a configuration that contains two distinct data identifiers, say  $d_1$  and  $d_2$ , that are both classes with the same label  $l$ . Thus, the formula  $t\text{-body}(d_1, d_2, l)$  is true. However, there is not an equivalent  $t\text{-sat}$  formula:  $t\text{-sat}(d_1, d_2, l)$  is false because  $d_1$  does not equal  $d_2$  (the  $C1=C2$  expression is false).

We show below that the two representations for constraints are equivalent, that is, constraints expressed in the abbreviated syntax and the evaluation of the stratified Datalog $\neg$  programs they generate. We start by showing that the generated stratified Datalog $\neg$  syntax is well-formed.

**Theorem 1 (Well-formed Datalog-program generation)** *Given a valid ULD constraint  $t$ , the Datalog program  $\mathcal{D}$  generated from the constraint is a well-formed stratified Datalog $\neg$  program.*

**Proof (Well-formed Datalog-program generation).** To show that  $\mathcal{D}$  can be stratified, notice that  $t$ 's constraint program is by definition a valid stratified Datalog $\neg$  program (because  $t\text{-body}$  and  $t\text{-sat}$  are not allowed as intensional predicate names for the IDB formulas  $r_1$  to  $r_m$ ). Assume  $P_1, P_2, \dots, P_i$  are used to partition the rules from  $t$ 's constraint program, where  $i$  is the number of needed partitions. The rule for  $t\text{-body}$  is placed in partition  $P_{i+1}$ , because  $b_c$  may contain negative IDB (constraint program) formulas. The rule for  $t\text{-sat}$  is also placed in  $P_{i+1}$  because it contains a positive (not a negative)  $t\text{-body}$  formula. While  $t\text{-sat}$  can contain negative formulas, they are restricted to extensional formulas only.

Finally,  $t\text{-sat}$  and  $t\text{-body}$  rules are valid Datalog $\neg$  rules as follows. By definition,  $b_c$  must be a valid Datalog $\neg$  rule body. If  $\text{var}(b_c)$  is empty,  $t\text{-body}$  is still well-formed because a constant term is added to the formula. For non-empty  $\text{var}(b_c)$ , every variable in the head of the  $t\text{-sat}$  rule also occurs in the body of the rule. The head of a constraint assertion may only contain a negative formula if each variable in the formula occurs in a positive formula in the head or body of the rule. Note that all variables that occur in the head of the assertion occur in the body of the  $t\text{-sat}$  rule. Therefore, the  $t\text{-sat}$  rule satisfies the Datalog $\neg$  restriction concerning negative formulas: For a given negative formula  $e_i$  in an

assertion, each variable of  $e_i$  either occurs in a positive formula  $e_j$  in the body of  $t\text{-sat}$ , or in the positive  $t\text{-body}$  formula in the body of  $t\text{-sat}$ . *QED*

Finally, we show that the two representations for constraints are semantically equivalent. That is, if  $t'$  is the constraint defined by the stratified Datalog $\neg$  program  $\mathcal{D}$  and  $t$  is a valid ULD constraint, we say that  $t$  and  $\mathcal{D}$  are *semantically equivalent* if for any configuration  $\mathcal{F}$  with corresponding database instance mapping  $I_{\mathcal{F}}$ ,  $I_{\mathcal{F}} \models t$  if and only if  $I_{\mathcal{F}} \models t'$ .

**Theorem 2 (Semantic equivalence of Datalog-program generation)** *Given any valid ULD constraint  $t$  and Datalog program  $\mathcal{D}$  generated from  $t$ ,  $\mathcal{D}$  and  $t$  are semantically equivalent.*

**Proof (Semantic equivalence of Datalog-program generation).** Assume that  $\mathcal{F}$  is a given configuration with database instance mapping  $I_{\mathcal{F}}$ . A variable assignment  $v_b$  makes  $b_c$  true if and only if  $v_b$  makes  $t\text{-body}$  true. For a given variable assignment  $v_b$  that makes  $b_c$  true,  $t\text{-sat}$  is true for  $v_b$  if and only if the expression  $e_1, e_2, \dots, e_n$  is true. It follows that  $I_{\mathcal{F}} \models t'$ , where  $t'$  represents the constraint imposed by  $\mathcal{D}$ , if only if for every  $v_b$  that makes  $b_c$  true,  $t\text{-sat}$  is also true. (Note that by definition of the rewriting,  $t'$  and  $t$  have the same  $b_c$  and  $e_1$  to  $e_n$ .) Thus,  $I_{\mathcal{F}} \models t'$  if and only if each  $e_i/v_b$  is true for  $1 \leq i \leq n$ , which is exactly equivalent to the conditions given by the ULD constraint definition for determining whether  $I_{\mathcal{F}} \models t$  is true. That is,  $e_i/v_b$  is true for  $1 \leq i \leq n$  if and only if the three conditions given by the ULD constraint definition hold. *QED*

The rest of this section gives examples of the ULD constraint language. The following constraint (Example 8) uses the *subclassClosure* (intensional) relation to ensure all RDF objects are instances of their corresponding classes (which may not be explicitly represented in an RDF source, but is usually inferred by an RDF processor). Note that the *subclassClosure* relation computes the transitive closure of the RDF subclass relationship.

**Example 8** (*RDF(S) subclass instance constraint*)

constraint obj1:

$$d\text{-inst}(R, C2) \Leftarrow c\text{-inst}(R, \text{resource}), d\text{-inst}(R, C1), \text{subclassClosure}(C1, C2).$$

$$\text{subclassClosure}(C1, C2) \Leftarrow C1=C2.$$

$$\begin{aligned} \text{subclassClosure}(C1, C3) \Leftarrow & c\text{-inst}(S, \text{subclass}), S.\text{hasSub}:C1, \\ & S.\text{hasSuper}:C2, \text{subclassClosure}(C2, C3). \end{aligned}$$

In general, constraints express information concerning configurations that are not directly expressible using the available structures of the ULD. Constraints can be used for a number of purposes such as validating that configurations (or the result of updating configurations) are correct, optimizing queries [2, 67], and transformation [68].

An important use of constraints for the ULD is to further define conformance relationships. To illustrate, Example 9 describes some of the restrictions imposed by conformance relationships in the relational model of Example 1. Note that in the relational model definition of Example 1, the conformance relationships only require tables and tuples to conform to relation types (through cardinality constraints) but they do not describe the structural implications of the conformance relationships. The constraints of Example 9 capture some of these structural requirements.

**Example 9** (*Relational conformance constraints*)

constraint rel-conf1:

$$d\text{-inst}(T, R) \Leftarrow c\text{-inst}(B, \text{table}), c\text{-inst}(R, \text{relation}), d\text{-inst}(B, R), T \in B.$$

constraint rel-conf2:

$$\begin{aligned} d\text{-inst}(V, D) \Leftarrow & c\text{-inst}(T, \text{tuple}), c\text{-inst}(R, \text{relation}), d\text{-inst}(T, R), R.\text{hasAtts}:AL, \\ & AL[I]=A, A.\text{hasDomain}:D, T[I]=V. \end{aligned}$$

constraint rel-conf3:

$$|T|=N \Leftarrow c\text{-inst}(T, \text{tuple}), c\text{-inst}(R, \text{relation}), d\text{-inst}(T, R), R.\text{hasAtts}:AL, |AL|=N.$$

The first constraint (**rel-conf1**) states that, for a table to conform to a relation, each tuple in the table must be a *d-inst* of the relation type. The second and third constraints define the conformance relationship between tuples and relations. The second constraint (**rel-conf2**) requires the *i*-th value in the tuple to be a *d-inst* of the domain of the relation's

$i$ -th attribute. The last constraint (**rel-conf3**) states that a tuple must have the same number of values as the number of attributes in the relation type.

Example 10 defines a conformance constraint for the RDF(S) data model of Example 3. Note that RDF(S) places few restrictions on conformance, for example, objects are not required to have any properties of their corresponding classes (that is, properties are optional). The constraint below, however, states that, if a triple's predicate is a defined property, the triple's subject must be a *d-inst* of the property's domain class, and the triple's object must be a *d-inst* of the property's range class. Note that this is an implicit conformance constraint, because there is not an explicit conformance definition between triples and properties.

**Example 10** (*RDF(S) conformance constraint*)

constraint **rdf-conf1**:

$$\begin{aligned} & \text{d-inst}(R_s, C_s), \text{d-inst}(R_o, C_o) \Leftarrow \\ & \text{c-inst}(T, \text{triple}), \text{c-inst}(P, \text{property}), T.\text{hasPred}:P, T.\text{hasSubj}=R_s, \\ & P.\text{hasDomain}=C_s, T.\text{hasObj}:R_o, P.\text{hasRange}:C_o. \end{aligned}$$

Not all constraints can be expressed using the ULD constraint language described in this section. For example, it is not possible to derive a disjunction, in which the intent of the constraint is to ensure that either  $e_1$  *or*  $e_2$  must hold. It is also not possible to express constraints that cannot be stratified. The ULD constraint language is slightly more expressive than stratified Datalog $\neg$  programs, because it permits existentially quantified and negative formulas in the head of an assertion. We show in Chapter 3 that the ULD constraint language is expressive enough to capture a wide range of data-model constraints.

## 2.6 A Partial ULD Axiomatization

The ULD provides uniform access to data model, schema, and data information. This flexibility allows generic restrictions to be placed on configurations using the ULD constraint language defined in the previous section. In other words, it is possible to define

general constraints (applicable to any data model) that must be satisfied by all valid configurations using the ULD constraint language. This section exploits this observation to further formalize the ULD framework. In particular, we provide an axiomatization for the ULD using the ULD constraint language. By *axiomatization*, we mean a set of statements (*axioms*) that must be true for all configurations. If an axiom is false for a configuration, the configuration is not considered valid. That is, the configuration is not a model of the axioms. Note that the axiomatization is independent of the syntax used to represent a configuration, and we assume that all configurations are syntactically valid. The axiomatization we present in this section is partial, and attempts to characterize the description of the ULD given previously in Sections 2.1, 2.2, and 2.3. There are true statements that can be made of valid configurations that are not included in the axiomatization. First, not all true statements can be expressed using the ULD constraint language. Examples include axioms for equality, collection arithmetic (such as asserting that the length of a collection is always the same as the the number of elements in the collection), and some closure axioms that require disjunction in the head of a rule. Second, we do not include all “trivial” axioms (mainly for conciseness), for example, that a domain constraint can only contain a value in the set  $\{1, +, ?, *\}$ .

We divide the axiomatization into two parts. First, we give a core set of axioms that are true regardless of the meta-data-model chosen. And second, we give the specific axioms for the ULD meta-data-model (with tuple, collection, union, and atomic structures). Although we do not consider alternative meta-data-models, it is possible to define different structures and constraints by choosing different axioms than those given for the ULD meta-data-model.<sup>6</sup>

When defining axioms, we use the conventions for variable names as shown in Figure 2.8. That is, we use variables  $S$  and  $T$  for construct types,  $P$  and  $Q$  for constructs, and  $X$  and  $Y$  for construct instances. When appropriate, we use the convention that  $P$  is a *ct-inst* of  $S$ ,  $Q$  is a *ct-inst* of  $T$ ,  $P$  can conform to  $Q$ ,  $X$  is a *c-inst* of  $P$ ,  $Y$  is a *c-inst* of  $Q$ , and  $X$  is a *d-inst* of  $Y$ . We also use the term **axiom** instead of **constraint** to define the

---

<sup>6</sup>In addition to defining an alternative meta-data-model, the syntax used to specify configurations would also need to be altered along with the mapping from the syntax to the corresponding database instance.




Figure 2.8: The naming conventions used in specifying ULD axioms.

rules in the theory. Semantically, constraints and axioms are identical, however, we use **constraint** to represent restrictions placed on specific data models and **axiom** to represent restrictions placed over all configurations regardless of the data model.

### 2.6.1 Definition of the Core ULD Theory

The set of axioms shown in Figure 2.9 and 2.10 are independent of any particular ULD meta-data-model. The axioms define the *core theory* of the ULD, which performs three primary functions: (1) it describes the dependency among construct types, constructs, and construct instances, (2) it partitions construct types, constructs, and construct instances into distinct sets, and (3) it articulates the relationship between conformance specifications (*conf* facts) and data instance-of relationships (*d-inst* facts). Note that we use the symbol \* to distinguish axioms that follow directly from the interpretation given in Section 2.4.2.

The first three axioms of Figure 2.9 describe the dependencies between construct types, constructs, and construct instances. The first axiom requires every construct identifier used in a configuration to be a *ct-inst* of some construct type. Recall that the  $\Leftarrow$  symbol denotes a constraint assertion in which existential variables are permitted in the head of the rule. Thus, the formula  $ct-inst(P, S)$  requires at least one  $S$  for each  $P$  in the configuration, where the equivalent first-order logic expression is:  $\forall X \forall P (c-inst(X, P) \Rightarrow \exists S ct-inst(P, S))$ . The second axiom requires the participating identifiers of every *d-inst* relationship to have constructs that occur together in a conformance definition. The third axiom requires conformance definitions to relate constructs (as opposed to construct types

- axiom 1: (*Constructs always have construct types.*)  
 $\text{ct-inst}(P, S) \Leftarrow \text{c-inst}(X, P).$
- axiom 2: (*Data instances always have conforming constructs.*)  
 $\text{c-inst}(X, P), \text{c-inst}(Y, Q), \text{conf}(P, Q, D, R) \Leftarrow \text{d-inst}(X, Y).$
- axiom 3\*: (*Conformance is defined between constructs.*)  
 $\text{ct-inst}(P, S), \text{ct-inst}(Q, T) \Leftarrow \text{conf}(P, Q, D, R).$
- axiom 4\*: (*Construct and construct types are disjoint.*)  
 $S \neq Q \Leftarrow \text{ct-inst}(P, S), \text{ct-inst}(Q, T).$
- axiom 5\*: (*Constructs and construct instances are disjoint.*)  
 $P \neq Y \Leftarrow \text{ct-inst}(P, S), \text{c-inst}(Y, Q).$
- axiom 6\*: (*Construct types and construct instances are disjoint.*)  
 $S \neq Y \Leftarrow \text{ct-inst}(P, S), \text{c-inst}(Y, Q).$
- axiom 7: (*Domain and Range constraints are unique among conforming constructs*)  
 $D1=D2, R1=R2 \Leftarrow \text{conf}(P, Q, D1, R1), \text{conf}(P, Q, D2, R2).$

Figure 2.9: The core ULD axioms for describing basic structure restrictions.

or construct instances).

Axioms 4 through 6 give pairwise restrictions on ULD meta-data-model identifiers. Axiom 4 requires the set of construct types and constructs to be disjoint. Axiom 5 requires the set of constructs and construct instances to be disjoint. And Axiom 6 requires the set of construct types and construct instances to be disjoint.

Axiom 7 requires unique conformance definitions, that is, if  $P$  can conform to  $Q$ , there is only one such definition for  $P$  and  $Q$  in a configuration. Note that conformance is not a symmetric relationship. However, it is possible to define separate conformance relationships such that  $P$  can conform to  $Q$  and  $Q$  can conform to  $P$  (but asserting  $P$  conforms to  $Q$  does not imply  $Q$  conforms to  $P$ ).

The remaining axioms of the core theory are shown in Figure 2.10 and define the domain and range cardinality restrictions on conformance. Recall that a conformance definition takes the form  $\text{conf}(P, Q, D, R)$ , which states that an instance  $X$  of  $P$  can conform to an instance  $Y$  of  $Q$  (that is, there can be a  $\text{d-inst}(X, Y)$  fact in the configuration). In

- axiom 8: (*Range constraint requiring at least one conformee.*)  
 $\text{c-inst}(Y, Q), \text{d-inst}(X, Y) \Leftarrow \text{conf}(P, Q, D, 1), \text{c-inst}(X, P).$
- axiom 9: (*Domain constraint requiring at least one conformer.*)  
 $\text{c-inst}(X, P), \text{d-inst}(X, Y) \Leftarrow \text{conf}(P, Q, 1, D), \text{c-inst}(Y, Q).$
- axiom 10: (*Range constraint requiring at least one conformee.*)  
 $\text{c-inst}(Y, Q), \text{d-inst}(X, Y) \Leftarrow \text{conf}(P, Q, D, +), \text{c-inst}(X, P).$
- axiom 11: (*Domain constraint requiring at least one conformer.*)  
 $\text{c-inst}(X, P), \text{d-inst}(X, Y) \Leftarrow \text{conf}(P, Q, +, D), \text{c-inst}(Y, Q).$
- axiom 12: (*Range constraint requiring at most one conformee.*)  
 $Y_1=Y_2 \Leftarrow \text{conf}(P, Q, D, 1), \text{c-inst}(X, P), \text{c-inst}(Y_1, Q), \text{c-inst}(Y_2, Q), \text{d-inst}(X, Y_1), \text{d-inst}(X, Y_2).$
- axiom 13: (*Domain constraint requiring at most one conformer.*)  
 $X_1=X_2 \Leftarrow \text{conf}(P, Q, 1, R), \text{c-inst}(Y, Q), \text{c-inst}(X_1, P), \text{c-inst}(X_2, P), \text{d-inst}(X_1, Y), \text{d-inst}(X_2, Y).$
- axiom 14: (*Range constraint requiring at most one conformee.*)  
 $Y_1=Y_2 \Leftarrow \text{conf}(P, Q, D, ?), \text{c-inst}(X, P), \text{c-inst}(Y_1, Q), \text{c-inst}(Y_2, Q), \text{d-inst}(X, Y_1), \text{d-inst}(X, Y_2).$
- axiom 15: (*Domain constraint requiring at most one conformer.*)  
 $X_1=X_2 \Leftarrow \text{conf}(P, Q, ?, R), \text{c-inst}(Y, Q), \text{c-inst}(X_1, P), \text{c-inst}(X_2, P), \text{d-inst}(X_1, Y), \text{d-inst}(X_2, Y).$

Figure 2.10: The core ULD axioms describing conformance cardinality restrictions.

addition,  $D$  represents a constraint on the domain (that is, the  $P$  side) of the relationship, and  $R$  represents a constraint on the range (that is, the  $Q$  side) of the relationship. For example, if the range cardinality restriction is 1, every instance  $X$  of  $P$  must conform to some instance  $Y$  of  $Q$  (that is,  $\forall X \exists Y (\text{d-inst}(X, Y))$ ). Similarly, if the domain cardinality restriction is 1, there must exist an instance  $X$  of  $P$  for every instance  $Y$  of  $Q$  such that  $X$  conforms to  $Y$  (that is,  $\forall Y \exists X (\text{d-inst}(X, Y))$ ).

We note that the core theory is less restrictive structurally than the ULD meta-data-model. In particular, the core theory permits a construct to be an instance of multiple construct types and to conform to multiple constructs.

axiom 16\*: (*Struct construct-type introduction.*)  
 $S = \text{struct-ct} \Leftarrow \text{ct-inst}(P, S), \neg S = \text{set-ct}, \neg S = \text{bag-ct}, \neg S = \text{list-ct}, \neg S = \text{atomic-ct}.$

axiom 17\*: (*Set construct-type introduction.*)  
 $S = \text{set-ct} \Leftarrow \text{ct-inst}(P, S), \neg S = \text{struct-ct}, \neg S = \text{bag-ct}, \neg S = \text{list-ct}, \neg S = \text{atomic-ct}.$

axiom 18\*: (*Bag construct-type introduction.*)  
 $S = \text{bag-ct} \Leftarrow \text{ct-inst}(P, S), \neg S = \text{struct-ct}, \neg S = \text{set-ct}, \neg S = \text{list-ct}, \neg S = \text{atomic-ct}.$

axiom 19\*: (*List construct-type introduction.*)  
 $S = \text{list-ct} \Leftarrow \text{ct-inst}(P, S), \neg S = \text{struct-ct}, \neg S = \text{set-ct}, \neg S = \text{bag-ct}, \neg S = \text{atomic-ct}.$

axiom 20\*: (*Atomic construct-type introduction.*)  
 $S = \text{atomic-ct} \Leftarrow \text{ct-inst}(P, S), \neg S = \text{struct-ct}, \neg S = \text{set-ct}, \neg S = \text{bag-ct}, \neg S = \text{list-ct}.$

axiom 21: (*A construct is an instance of at most one construct type.*)  
 $S_1 = S_2 \Leftarrow \text{ct-inst}(P, S_1), \text{ct-inst}(P, S_2).$

axiom 22: (*A construct conforms to at most one construct.*)  
 $Q_1 = Q_2 \Leftarrow \text{conf}(P, Q_1, D_1, R_1), \text{conf}(P, Q_2, D_2, R_2).$

Figure 2.11: The ULD meta-data-model axioms that define the allowable construct types and two structural restrictions.

## 2.6.2 Definition of the ULD Meta-Data-Model Theory

The ULD meta-data-model axioms are given in Figures 2.11–2.17. The meta-data-model axioms extend the core theory by defining the allowable construct types and the corresponding representations for basic structures.

The first five axioms (16–22) of Figure 2.11 restrict the set of construct-type constants to *struct-ct*, *set-ct*, *bag-ct*, *list-ct*, and *atomic-ct*. Thus, any configuration that uses a different construct type than those defined is not considered valid. Taken together, Axioms 16–20 are equivalent to the first-order logic expression  $\forall P \forall S (\text{ct-inst}(P, S) \Rightarrow S = \text{struct-ct} \vee S = \text{set-ct} \vee S = \text{bag-ct} \vee S = \text{list-ct} \vee S = \text{atomic-ct} \vee S = \text{union-ct})$ , however, disjunction constraints are not allowed directly. The last two axioms (20 and 21) define structural restrictions on configurations. In particular, Axiom 20 restricts a construct from being a *ct-inst* of more than one construct type, and Axiom 22 refines Axiom 7 of the core theory by restricting a construct to conform to at most one other construct.

Figure 2.12 defines axioms for tuple structures (that is, *struct-ct* basic structures). The

- axiom 23\*: (*Component definitions exist only for struct constructs.*)  
 $\text{ct-inst}(P, \text{struct-ct}), \text{ct-inst}(Q, T) \Leftarrow P.\text{Sel} \rightarrow Q.$
- axiom 24\*: (*Component expressions exist only for struct construct instances.*)  
 $\text{c-inst}(X, P), \text{ct-inst}(P, \text{struct-ct}), \text{c-inst}(Y, Q) \Leftarrow X.\text{Sel}:Y.$
- axiom 25: (*Every struct construct has at least one component definition.*)  
 $P.\text{Sel} \rightarrow Q \Leftarrow \text{ct-inst}(P, \text{struct-ct}).$
- axiom 26: (*Struct instances always have a component definition for each component.*)  
 $X.\text{Sel}:Y, \text{c-inst}(Y, Q) \Leftarrow \text{ct-inst}(P, \text{struct-ct}), P.\text{Sel} \rightarrow Q, \text{c-inst}(X, P).$
- axiom 27: (*Struct instances always have components for component definitions.*)  
 $P.\text{Sel} \rightarrow Q, \text{c-inst}(Y, Q) \Leftarrow \text{ct-inst}(P, \text{struct-ct}), \text{c-inst}(X, P), X.\text{Sel}:Y.$
- axiom 28: (*Struct constructs have unique selectors.*)  
 $Q_1 = Q_2 \Leftarrow \text{ct-inst}(P, \text{struct-ct}), P.\text{Sel} \rightarrow Q_1, P.\text{Sel} \rightarrow Q_2.$
- axiom 29: (*Struct instances have unique selectors.*)  
 $Y_1 = Y_2 \Leftarrow \text{ct-inst}(P, \text{struct-ct}), \text{c-inst}(X, P), P.\text{Sel}:Y_1, P.\text{Sel}:Y_2.$

Figure 2.12: The axioms for tuple structures in the ULD meta-data-model.

first two axioms (23 and 24) define component expressions. For example, Axiom 23 requires  $P$  to be a *struct-ct* construct and  $Q$  to be a construct in a component definition  $P.\text{Sel} \rightarrow Q$ . Axiom 25 requires *struct-ct* constructs to have at least one component definition. Axioms 26 and 27 require each construct to have exactly one component for each corresponding component definition. And Axioms 28 and 29 require *struct-ct* constructs and construct instances to have unique component selectors.

Figure 2.13 defines axioms for collections (that is, *set-ct*, *bag-ct*, and *list-ct* basic structures). Axiom 30 requires that membership occurs between construct instances and that membership formulas are applied only to collections. Axiom 31 defines the same restrictions as Axiom 30 for collection-length formulas.

Figures 2.14–2.16 axiomatize the set, bag, and list structures, respectively. Axioms 32, 37, and 44 require collection definitions to occur only for collection constructs. Axioms 33, 38, and 45 require every collection construct to have a collection definition. Axioms 34, 39, and 46 require a collection construct to have at most one collection definition. Axioms 35, 40, and 47 require collection instances to consist of only valid members, that is, members

axiom 30\*: (*The membership predicate only applies to collection construct instances.*)  
 $\text{c-inst}(X, P), \text{c-inst}(Y, Q), \neg \text{ct-inst}(P, \text{struct-ct}), \neg \text{ct-inst}(P, \text{atomic-ct}) \Leftarrow Y \in X.$

axiom 31\*: (*Only collections have an element count.*)  
 $\text{c-inst}(X, P), \neg \text{ct-inst}(P, \text{struct-ct}), \neg \text{ct-inst}(P, \text{atomic-ct}) \Leftarrow |X|=C.$

Figure 2.13: Two general axioms for collection structures in the ULD meta-data-model.

axiom 32\*: (*Set definitions are between a set construct and another construct.*)  
 $\text{ct-inst}(P, \text{set-ct}), \text{ct-inst}(Q, T) \Leftarrow \text{setof}(P, Q).$

axiom 33\*: (*Set constructs always have a set definition.*)  
 $\text{setof}(P, Q) \Leftarrow \text{ct-inst}(P, \text{set-ct}).$

axiom 34: (*Set constructs have at most one set definition.*)  
 $Q_1=Q_2 \Leftarrow \text{ct-inst}(P, \text{set-ct}), \text{setof}(P, Q_1), \text{setof}(P, Q_2).$

axiom 35: (*Set instances always have valid members.*)  
 $\text{c-inst}(Y, Q) \Leftarrow \text{ct-inst}(P, \text{set-ct}), \text{setof}(P, Q), \text{c-inst}(X, P), Y \in X.$

axiom 36\*: (*Set instances always have a length.*)  
 $|X|=C \Leftarrow \text{ct-inst}(P, \text{set-ct}), \text{c-inst}(X, P).$

Figure 2.14: The axioms for set structures in the ULD meta-data-model.

that are instances of the construct mentioned in the collection definition. And axioms 36, 41, and 48 require collection instances to have a cardinality.

Axioms 42 and 43 of Figure 2.15 define the duplicate-element predicate for bags. Recall that for the bag  $x$ , construct-instance  $y$ , and integer  $c$ , the expression  $y \in\in x = c$  states that  $y$  occurs  $c$  times as a member of  $x$ . Axiom 42 requires each member of a bag to have an associated duplicate-element formula. And Axiom 43 limits duplicate-element formulas to bag construct instances.

Axioms 49 and 50 of Figure 2.16 define the list-index predicate. Axiom 49 requires each member in a list to have an associated index in the list. And Axiom 50 limits list-index expressions to members of list construct instances.

Figure 2.17 gives the axioms for unions (that is, *union-ct* constructs). Axioms 51

- axiom 37\*: (*Bag definitions are between a bag construct and another construct.*)  
 $\text{ct-inst}(P, \text{bag-ct}), \text{ct-inst}(Q, T) \Leftarrow \text{bagof}(P, Q).$
- axiom 38\*: (*Bag constructs always have a bag definition.*)  
 $\text{bagof}(P, Q) \Leftarrow \text{ct-inst}(P, \text{bag-ct}).$
- axiom 39: (*Bag constructs have at most one bag definition.*)  
 $Q_1=Q_2 \Leftarrow \text{ct-inst}(P, \text{bag-ct}), \text{bagof}(P, Q_1), \text{bagof}(P, Q_2).$
- axiom 40\*: (*Bag instances always have valid members.*)  
 $\text{c-inst}(Y, Q) \Leftarrow \text{ct-inst}(P, \text{bag-ct}), \text{bagof}(P, Q), \text{c-inst}(X, P), Y \in X.$
- axiom 41\*: (*Bag instances always have a length.*)  
 $|X|=C \Leftarrow \text{ct-inst}(P, \text{bag-ct}), \text{c-inst}(X, P).$
- axiom 42\*: (*Bag instances always have a duplicate-element expression for each member.*)  
 $Y \in X=C \Leftarrow \text{ct-inst}(P, \text{bag-ct}), \text{c-inst}(X, P), Y \in X.$
- axiom 43\*: (*Only bags have duplicate-element expressions.*)  
 $\text{c-inst}(X, P), \text{ct-inst}(P, \text{bag-ct}), \text{c-inst}(Y, Q) \Leftarrow Y \in X=C.$

Figure 2.15: The axioms for bag structures in the ULD meta-data-model.

and 52 require that every union definition is for a union construct (Axiom 51) and union constructs have at least one union definition (Axiom 52). Axiom 53 requires each instance of a union construct to also be an instance of another, non-union construct. Axiom 54 requires each instance of a construct that participates in a union construct to be an instance of the union construct. Axiom 55 disallows union constructs from conforming to any other construct. And Axiom 56 restricts union definitions from being acyclic.

Besides serving as a formal description of the ULD and the ULD meta-data-model, the axiomatization described in this section can be used to (partially) validate a configuration directly. That is, by using the ULD constraint language to express axioms, we could execute these axioms using a Datalog or Prolog system to check that a configuration is valid.

In addition, we believe the axiomatization can be used to infer new facts concerning configurations and to derive additional statements, for further reasoning. One use of the axiomatization is to automate the population of a configuration, for example, when we

- axiom 44\*: (*List definitions are between a list construct and another construct.*)  
 $\text{ct-inst}(P, \text{list-ct}), \text{ct-inst}(Q, T) \Leftarrow \text{listof}(P, Q).$
- axiom 45\*: (*List constructs always have list definitions.*)  
 $\text{listof}(P, Q) \Leftarrow \text{ct-inst}(P, \text{list-ct}).$
- axiom 46: (*List constructs have at most one definition.*)  
 $Q_1=Q_2 \Leftarrow \text{ct-inst}(P, \text{list-ct}), \text{listof}(P, Q_1), \text{listof}(P, Q_2).$
- axiom 47\*: (*List instances have only valid members.*)  
 $\text{c-inst}(Y, Q) \Leftarrow \text{ct-inst}(P, \text{list-ct}), \text{listof}(P, Q), \text{c-inst}(X, P), Y \in X.$
- axiom 48\*: (*List instances always have a length.*)  
 $|X|=C \Leftarrow \text{ct-inst}(P, \text{list-ct}), \text{c-inst}(X, P).$
- axiom 49\*: (*List instances always have an index expression for each member.*)  
 $X[I]=Y \Leftarrow \text{ct-inst}(P, \text{list-ct}), \text{c-inst}(X, P), Y \in X.$
- axiom 50\*: (*Only lists have indexes.*)  
 $\text{c-inst}(X, P), \text{ct-inst}(P, \text{list-ct}), \text{c-inst}(Y, Q) \Leftarrow X[I]=Y.$

Figure 2.16: The axioms for list structures in the ULD meta-data-model.

wish to transform one configuration into another configuration or an information source into its corresponding configuration. We describe this use of constraints in more detail in Chapter 4.

As a simple example, assume we have defined a rule to convert existing information into an RDF resource (using the RDF(S) data model defined in Example 3). However, we would like to know if an instance of an RDF resource must be an instance of any other construct (so our program can define a valid configuration). This question can be answered using Axiom 54, which states that any instance of a construct that participates in a union definition is also an instance of the union construct. Therefore, we can use Axiom 54 to prove that every **resource** construct instance must also be an **objVal** constant instance.

As another example, using Axioms 2 and 22 we can derive the following true statement concerning configurations.

$$\text{c-inst}(P, Q) \Leftarrow \text{d-inst}(X, Y), \text{c-inst}(X, P), \text{conf}(P, Q, D, R).$$

This axiom allows us to infer the construct of an instance that participates in a *d-inst*

axiom 51\*: (*Union definitions exist only for union constructs.*)  
 $\text{ct-inst}(P, \text{union-ct}), \text{ct-inst}(Q, T) \Leftarrow \text{unionof}(P, Q).$

axiom 52: (*Union constructs are composed of at least two distinct constructs.*)  
 $\text{unionof}(P, Q_1), \text{unionof}(P, Q_2), \neg Q_1=Q_2, \Leftarrow \text{ct-inst}(P, \text{union-ct}).$

axiom 53: (*Union constructs do not have direct instances.*)  
 $\text{c-inst}(X, Q), \neg \text{ct-inst}(Q, \text{union-ct}) \Leftarrow \text{ct-inst}(P, \text{union-ct}), \text{c-inst}(X, P).$

axiom 54\*: (*Instances of constructs grouped by union are instances of the union.*)  
 $\text{c-inst}(Y, P) \Leftarrow \text{ct-inst}(P, \text{union-ct}), \text{unionof}(P, Q), \text{c-inst}(Y, Q).$

axiom 55: (*Union constructs do not conform to other constructs.*)  
 $\neg \text{conf}(P, Q, D, R) \Leftarrow \text{ct-inst}(P, \text{union-ct}), \text{ct-inst}(Q, T), \text{conf}(P', Q', D, R).$

axiom 56: (*Union constructs are acyclic.*)  
 $P \neq P' \Leftarrow \text{c-inst}(P, \text{union-ct}), \text{unionConnected}(P, P').$   
 $\text{unionConnected}(P_1, P_2) \Leftarrow \text{c-inst}(P_1, \text{union-ct}), \text{unionof}(P_1, P_2).$   
 $\text{unionConnected}(P_1, P_3) \Leftarrow \text{c-inst}(P_1, \text{union-ct}), \text{unionof}(P_1, P_2), \text{unionConnected}(P_2, P_3).$

Figure 2.17: The axioms for union definitions in the ULD meta-data-model.

relationship. For example, assuming we have access to the RDF(S) data-model definition, we can use this new axiom with Axiom 27 to determine that the following query is not well-formed (we discuss these issues in more detail in Chapter 4).

$$\text{ans}(R) \Leftarrow \text{d-inst}(X, Y), \text{c-inst}(X, \text{resource}), Y.\text{hasRange}:R.$$

In particular, because we know (from the definition of the data model) that a resource can conform to a class (therefore,  $Y$  should be an instance of a class according to our new axiom) and classes do not have range selectors (from Axiom 27), we can infer that the above query is not valid.

## 2.7 Summary

In this chapter, we defined the ULD and compared it to typical meta-data-model architectures. The ULD offers a flat representation for storing data model constructs, schema, and data. The ULD can also be used to represent data models that permit optional, partial, and multiple levels of schema. We gave a syntax and first-order logic interpretation

for representing configurations using the ULD. We introduced the ULD query language and constraint language. The query language is based on stratified Datalog $\neg$  and the constraint language extends the query language to permit existential restrictions. And finally, we used the constraint language to specify a partial axiomatization of the ULD.

In the next chapter, we use the ULD to define a wide variety of data models, including their constructs and constraints. The data models chosen demonstrate the power of the ULD to model various styles of schema conformance. We identify and give a set of shorthand notations for expressing common data-model constraints. And finally, we describe a set of properties that can be used to characterize data models in terms of their conformance requirements.

# Chapter 3

## Modeling Data Models

In this chapter, we use the ULD to give complete descriptions of a number of existing data models. Our goal is to support our claim that the ULD can represent a wide variety of data models. We consider both constructs and constraints of data models, including conformance definitions. To help specify structural constraints for defining data models, we introduce *constraint macros* in Section 3.1. A constraint macro is a short-hand notation for specifying commonly-used constraints over constructs. Section 3.2 presents ULD representations for the relational [33], E-R [27], and object-oriented [48] data models, which we consider traditional database representations. In Section 3.2, we extend the definition of the relational data model given in Chapter 2 with constraints. Section 3.3 presents ULD representations for the XML [22] and RDF [47] data models, which are primarily used for storing structured information on the Web. We consider XML with DTDs [22], which provide a schema language for XML, and RDF with RDF Schema [23], which serves as the base schema language for RDF, for example, OWL Lite, OWL DL, and OWL Full [56] are layered on top of RDF Schema. We extend the XML and RDF definitions given in Chapter 2 with constraints and additional constructs. In Section 3.4, we present the Topic Map [13] and VIKI [55] hypertext data models. The Topic Map data model is similar to RDF in that it can be used to annotate Web-based resources (such as Web-pages), however it consists of fundamentally different structuring primitives more closely related to structures found in hypertext systems than in semantic networks. The VIKI data model is a structured hypertext data model that is inherently “schema-later” in that tools based on VIKI allow users to create structured, un-typed data, but suggest possible schema types when appropriate. Finally, Section 3.5 introduces a scheme for classifying properties of

conformance, which we use to summarize the data models presented in this chapter.

There are a number of ways that constraints can be exploited for data models, and we briefly mention some uses in this chapter (we further discuss uses of constraints in Chapter 4). For a number of data models, constraints can be used not only to clarify and ensure valid configurations, but also to perform basic reasoning services over structures, such as type inference given a schema. In addition, because constraints can provide specification of conformance relationships, they can help further classify data models in terms of their allowed schema-instance arrangements.

### 3.1 Expressing Common Constraints

We use constraint macros to easily express common structural constraints for data models. Macros are typically more concise for expressing constraints than the constraint language introduced in Chapter 2, but they are not as general: Constraint macros can only be used for certain classes of constraints. (Note that one could define additional forms of macros, not described here, as needed.) We consider macros for expressing uniqueness constraints, which are similar to keys in a relational database, basic collection constraints, and inclusion constraints, which are similar to foreign keys in the relational data model. Each macro described in this section produces a corresponding expression in the constraint language defined in Chapter 2.

#### 3.1.1 Uniqueness Constraints

Given a construct and a set of non-empty component selectors, a *uniqueness constraint* requires each instance of the construct in a configuration to have unique values for the corresponding selectors. The syntax for a uniqueness constraint macro takes the following form, where  $c$  is the name (identifier) of a construct,  $s_1$  to  $s_n$  are component selectors of  $c$ ,  $n \geq 1$ , and  $t$  is the name of the constraint.

constraint  $t$ : (  $s_1, s_2, \dots, s_n$  keyOf  $c$  ).

The expression requires every instance of  $c$  to have unique values for  $s_1$  to  $s_n$  in a configuration. For example, the following constraint (taken from Chapter 2) expresses a typical

uniqueness restriction using the ULD constraint language. The constraint requires all RDF classes in a configuration to have a unique label value.

constraint key1:

$$C1=C2 \Leftarrow \text{c-inst}(C1, \text{class}), \text{c-inst}(C2, \text{class}), C1.\text{hasLabel:L}, C2.\text{hasLabel:L}.$$

This constraint can be equivalently expressed using the following uniqueness constraint macro.

constraint key1: ( hasLabel keyOf class ).

When a uniqueness constraint is expressed against a tuple construct (as in the previous example), the restriction produces the following assertion expressed in the ULD constraint language, where  $c$  is a construct,  $s_1$  to  $s_n$  are the given component selectors of  $c$ , and  $n \geq 1$ .

$$\begin{aligned} X_1=X_2 &\Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), \\ X_1.s_1:Y_1, X_1.s_2:Y_2, \dots, X_1.s_n:Y_n, \\ X_2.s_1:Y_1, X_2.s_2:Y_2, \dots, X_2.s_n:Y_n. \end{aligned}$$

If a uniqueness constraint requires every instance of a tuple construct to be unique in a configuration, we say the constraint is *configuration based*, that is, the constraint limits all construct instances within the configuration. In our previous example with RDF class labels, each instance of the class construct must contain a unique label value in the configuration. Thus, the constraint is configuration-based as are all uniqueness constraints expressed directly against tuple constructs.

Alternatively, a *collection-based* uniqueness constraint restricts tuple instances only within a collection as opposed to within the entire configuration. For a collection construct  $c_1$  defined over a tuple construct  $c_2$ , a collection-based uniqueness constraint restricts the members of  $c_1$  to those instances of  $c_2$  that have unique selector values. The macro syntax for defining collection-based uniqueness constraints is identical to the syntax for expressing configuration-based uniqueness constraints. (We use the type of construct being restricted to distinguish between collection-based and configuration-based uniqueness constraints.) For instance, the following collection-based uniqueness constraint requires

attribute lists (that is, instances of the **attList** construct) in the relational data model to contain uniquely named **attribute** instances. Thus, **attList** collections may not contain **attribute** items with the same name.

constraint key2: ( hasName keyOf attList )

A constraint expressed using a collection-based uniqueness macro produces the following ULD constraint language assertion, where  $c$  is the collection construct,  $s_1$  to  $s_n$  are the associated component selectors, and  $n \geq 1$ .

$$\begin{aligned} X_1=X_2 \Leftarrow & \text{c-inst}(X, c), X_1 \in X, X_2 \in X, \\ & X_1.s_1:Y_1, X_1.s_2:Y_2, \dots, X_1.s_n:Y_n, \\ & X_2.s_1:Y_1, X_2.s_2:Y_2, \dots, X_2.s_n:Y_n. \end{aligned}$$

Certain collection-based uniqueness constraints can be inferred from configuration-based uniqueness constraints. Namely, if a set construct  $c_1$  is defined over  $c_2$  and  $c_2$  has a uniqueness constraint over selectors  $s_1$  to  $s_n$ , a collection-based uniqueness constraint also implicitly exists for  $c_1$  over selectors  $s_1$  to  $s_n$ . In addition, constraints similar to relational superkeys [71] can be derived from uniqueness restrictions. For a construct with selectors  $S$  and a uniqueness constraint over selectors  $S_k$  for  $S_k \subseteq S$ , every set of selectors  $S_k \cup \{s\}$  for one or more  $s \in S - S_k$  is also a valid, although more general, uniqueness constraint for the construct.

Finally, we allow configuration-based uniqueness constraints to be defined over union constructs in addition to tuple constructs. For example, consider a union construct  $c$  that groups tuple constructs  $c_1$  and  $c_2$ , and both  $c_1$  and  $c_2$  share a component selector  $s$ . Defining  $s$  as a key for  $c$ , for example, using the uniqueness constraint ( $s$  keyOf  $c$ ), restricts the set of instances of  $c_1$  and  $c_2$  in a configuration such that each instance has a unique value for  $s$ . Note that for the case of union constructs, the same generated ULD constraint (above) still correctly defines uniqueness.

### 3.1.2 Collection Constraints

Here we describe two additional collection-based constraint macros. The first is called the *non-empty collection* constraint macro, which asserts that every instance of a given

collection construct in a configuration contains at least one item. The syntax and corresponding ULD constraint-language assertion for the non-empty collection macro is given below, where  $c$  is a collection construct and  $t$  is the name of the constraint.

constraint  $t$ : (  $c$  isNonEmpty )

$Y \in X \Leftarrow \text{c-inst}(X, c).$

The second macro is called the *distinct-collection* constraint and states that every instance of a given collection construct contains unique values. For example, given a list construct  $c$  with a distinct-collection constraint, each instance of  $c$  in a configuration must contain a unique collection of identifiers. The macro takes the following form, where  $c$  is a collection construct and  $t$  is the name of the constraint.

constraint  $t$ : (  $c$  isDistinct )

If  $c$  is a set construct, the following ULD constraint-language assertion is generated from the distinct-collection constraint macro.

$X_1 = X_2 \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), \neg \text{hasDifferentSetItem}(X_1, X_2).$

$\text{hasDifferentSetItem}(X_1, X_2) \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), Y \in X_1, \neg Y \in X_2.$

$\text{hasDifferentSetItem}(X_1, X_2) \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), Y \in X_2, \neg Y \in X_1.$

Similarly, if  $c$  is a bag construct, the following ULD constraint-language assertion is generated. Recall that the expression  $y \in\in x = n$  states that  $y$  is in the bag  $x$ ,  $n$  times.

$X_1 = X_2 \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), \neg \text{hasDifferentBagItem}(X_1, X_2).$

$\text{hasDifferentBagItem}(X_1, X_2) \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), Y \in X_1, \neg Y \in X_2.$

$\text{hasDifferentBagItem}(X_1, X_2) \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), Y \in X_2, \neg Y \in X_1.$

$\text{hasDifferentBagItem}(X_1, X_2) \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), Y \in X_1, Y \in X_2,$

$Y \in\in X_1 = N, Y \in\in X_2 \neq N.$

Finally, for a list construct  $c$ , the following ULD constraint-language assertion is generated from the distinct-collection macro.

$$X_1=X_2 \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), \neg\text{hasDifferentListItem}(X_1, X_2).$$

$$\text{hasDifferentListItem}(X_1, X_2) \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), X_1[I]=Y, X_2[I]\neq Y.$$

$$\text{hasDifferentListItem}(X_1, X_2) \Leftarrow \text{c-inst}(X_1, c), \text{c-inst}(X_2, c), X_2[I]=Y, X_1[I]\neq Y.$$

### 3.1.3 Inclusion Constraints

Inclusion constraint macros are used to ensure that certain construct instances always participate, that is, are always *included*, in some other construct instance. Here we consider macros for membership, joined-membership, joined-subset, and component inclusion constraints.

An example of an inclusion constraint occurs in the relational model, in which every tuple must belong to a table. In other words, it is not possible for a tuple to exist independently of a table. Because tuples are members of table collections, we call the restriction a *membership* inclusion constraint. The macro for a membership inclusion constraint takes the following form, where  $c_1$  is a construct,  $c_2$  is a collection construct defined over  $c_1$ , and  $t$  is the name of the constraint.

$$\text{constraint } t: ( c_1 \text{ always}(\in) c_2 ).$$

The constraint states that every instance of  $c_1$  in a configuration is always a member of an instance of  $c_2$ . For example, we would use the macro (tuple **always**( $\in$ ) table) to express the previous relational inclusion constraint. The following ULD constraint-language assertion is produced by a membership inclusion macro.

$$\text{c-inst}(Y, c_2), X\in Y \Leftarrow \text{c-inst}(X, c_1).$$

Often, it is useful to define inclusion constraints in the context of other constructs. For example, in the relational model, every attribute must be associated with a particular relation. This restriction can be captured using the following inclusion macro, where

the `relation` construct has the selector definition `hasAtts` that leads to the `attList` construct, which is defined as a list of attributes.

constraint inc-1: ( attribute `always`( $\in$ ) `relation.hasAtts` ).

The constraint requires that each attribute is a member of at least one relation's attribute list. As shown, we allow inclusion constraints to contain simple path expressions. A simple path expression has the form  $c.s_1.s_2\dots s_n$ , where  $c$  is a tuple construct (or a union of tuple constructs), and  $s_1$  to  $s_n$  are tuple selectors. We call  $c$  the *base* construct of the path expression. The following expression represents the general form for membership inclusion macros with simple path expressions, where  $c_1$  and  $c_2$  are base constructs,  $m \geq 0$ ,  $n \geq 0$ , and  $t$  is the name of the construct.

constraint  $t$ : (  $c_1.s_{11}.s_{12}\dots s_{1n}$  `always`( $\in$ )  $c_2.s_{21}.s_{22}\dots s_{2m}$  ).

The following ULD constraint-language assertion is produced by the general form of the membership inclusion macro.

$$\begin{aligned} &\text{c-inst}(Y, c_2), Y.s_{21}:Y_1, Y_1.s_{22}:Y_2, \dots, Y_{m-1}.s_{2m}:Y_m, X_n \in Y_m \Leftarrow \\ &\text{c-inst}(X, c_1), X.s_{11}:X_1, X_1.s_{12}:X_2, \dots, X_{n-1}.s_{1n}:X_n. \end{aligned}$$

To illustrate, the following ULD constraint is generated from the previous relational macro (`inc-1`).

$$\text{c-inst}(Y, \text{relation}), Y.\text{hasAtts}:Y_1, X \in Y_1 \Leftarrow \text{c-inst}(X, \text{attribute}).$$

A *joined* inclusion-constraint macro is a special case of an inclusion-constraint macro in which the base construct  $c_1$  is equal to the base construct  $c_2$  (that is,  $c_1$  and  $c_2$  are the same identifier). In a joined inclusion-constraint macro, we treat each side of the constraint as though it denotes the same base-construct instance. For example, for any instance  $x$  of construct  $c$  with collection instance  $y$  such that  $x.s:y$ , the joined inclusion-constraint macro ( $c$  `always` $\bowtie$ ( $\in$ )  $c.s$ ) requires  $x$  to be a member of  $y$ . (Note that we use `always` $\bowtie$  to distinguish a joined inclusion-constraint macro from a normal inclusion-constraint macro.)

The following expression represents the general form for joined membership inclusion-constraint macros, where  $c$  is the base construct,  $m \geq 0$ ,  $n \geq 0$ , and  $t$  is the name of the constraint.

$$\text{constraint } t: ( c.s_{11}.s_{12}...s_{1n} \text{ always}\bowtie(\in) c.s_{21}.s_{22}...s_{2m} ).$$

The following ULD constraint-language assertion is generated for joined membership inclusion-constraint macros, where  $c$  denotes the base construct.

$$\begin{aligned} X.s_{21}:Y_1, Y_1.s_{22}:Y_2, \dots, Y_{m-1}.s_{2m}:Y_m, X_n \in Y_m \Leftarrow \\ \text{c-inst}(X,c), X.s_{11}:X_1, X_1.s_{12}:X_2, \dots, X_{n-1}.s_{1n}:X_n. \end{aligned}$$

For example, consider the following joined membership-inclusion constraint for the E-R data model (see Section 3.2.3). A relationship type in the E-R data model contains a list of roles, which can optionally be named. The `roleName` construct assigns a name to a particular role for a relationship type. The constraint below requires the role being named to be a valid role for the relationship type.

$$\text{constraint inc-2: ( roleName.forRole always}\bowtie(\in) \text{ roleName.forType.hasRoles } ).$$

The following ULD constraint is generated from the macro.

$$X.\text{forType}:Y, Y.\text{hasRoles}:Y_1, X_1 \in Y_1 \Leftarrow \text{c-inst}(X, \text{roleName}), X.\text{forRole}:X_1.$$

We consider *subset* and *component* inclusion constraints in addition to membership inclusion constraints. A subset inclusion constraint occurs between two collection constructs<sup>1</sup> and requires a subset relationship between their instances. Here, we only consider joined subset inclusion-constraint macros, which take the following form, where  $c$  is the base construct,  $n \geq 0$ ,  $m \geq 0$ , and  $t$  is the name of the constraint. (We discuss general subset inclusion constraints below.)

$$\text{constraint } t: ( c.s_{11}.s_{12}...s_{1n} \text{ always}\bowtie(\subseteq) c.s_{21}.s_{22}...s_{2m} ).$$

---

<sup>1</sup>The subset relation is applied to lists and bags as if they were sets, that is, bags and lists are viewed as sets when testing for a subset relationship.

The following constraint is generated for each joined subset-inclusion macro for base construct  $c$ .

$$\begin{aligned} X.s_{21}:Y_1, Y_1.s_{22}:Y_2, \dots, Y_{m-1}.s_{2m}:Y_m, Z \in Y_m \Leftarrow \\ c\text{-inst}(X, c), X.s_{11}:X_1, X_1.s_{12}:X_2, \dots, X_{n-1}.s_{1n}:X_n, Z \in X_n. \end{aligned}$$

The next relational data-model constraint is an example of a joined subset-inclusion macro. A primary-key construct consists of a list of attributes and a relation (which is the relation that the primary key restricts). The constraint below states that the primary-key attributes are a subset of the attributes of the relation that the primary key is defined for.

$$\text{constraint inc-3: ( pKey.keyAtts always} \bowtie (\subseteq) \text{ pKey.forRel.hasAtts )}$$

The following ULD constraint is generated from the macro.

$$X.\text{forRel}:Y_1, Y_1.\text{hasAtts}:Y_2, Z \in Y_2 \Leftarrow c\text{-inst}(X, \text{pKey}), X.\text{keyAtts}:X_1, Z \in X_1.$$

A component inclusion-constraint macro requires every instance  $x$  of a construct  $c_1$  (or an instance of the path started by  $c_1$ ) to participate as a value of the selector  $s$  of some instance  $y$  of  $c_2$  (or an instance of the path started by  $c_2$ ). That is, in the case where no additional paths are given for  $c_1$  or  $c_2$ ,  $y.s:x$  must be true. The component inclusion macro takes the following form, where  $c_1$  and  $c_2$  are constructs,  $n \geq 0$ ,  $m \geq 0$ , and  $t$  is the name of the constraint.

$$\text{constraint } t: ( c_1.s_{11}.s_{12} \dots s_{1n} \text{ always}(=) c_2.s_{21}.s_{22} \dots s_{2m}.s ).$$

The following ULD constraint-language assertion is produced by a component inclusion macro for base constructs  $c_1$  and  $c_2$ .

$$\begin{aligned} c\text{-inst}(Y, c_2), Y.s_{21}:Y_1, Y_1.s_{22}:Y_2, \dots, Y_{m-1}.s_{2m}:Y_m, Y_m.s:X_n \Leftarrow \\ c\text{-inst}(X, c_1), X.s_{11}:X_1, X_1.s_{12}:X_2, \dots, X_{n-1}.s_{1n}:X_n. \end{aligned}$$

As an example, the following component inclusion macro requires all attList constructs in the relational model to occur in a relation.

constraint inc-4: (attList always(=) relation.hasAtts )

The following ULD constraint is generated from the macro.

c-inst(Y, relation), Y.hasAtts:X  $\Leftarrow$  c-inst(X, attList).

We do not consider joined component inclusion-constraint macros because none of the data models described in this chapter require the constraint. However, generating a ULD constraint assertion for a joined component inclusion constraint is straightforward.

Finally, we note the following two properties of inclusion constraint macros. First, for any construct  $c$ , the macro  $(c.s_1...s_n \text{ always}\bowtie(\subseteq) c.s_1...s_n)$  for  $n \geq 0$  is guaranteed to be satisfied in every configuration (assuming the macro is well-formed). Second, certain inclusion constraints are implied by other, more general inclusion constraints, which we list below. We use  $p$  and  $p'$  to denote arbitrary path expressions. In addition, we use the shorthand notation  $p \rightsquigarrow c$  (that is,  $p$  “leads to”  $c$ ) if  $p$  is a path expression  $c_1.s_1.s_2...s_n$  such that  $c_1.s_1 \rightarrow c_2$ ,  $c_2.s_2 \rightarrow c_3$ , ..., and  $c_{n-1}.s_n \rightarrow c$ , for  $n \geq 0$ . Note that, if  $n = 0$  and  $p = c$ , then trivially  $c \rightsquigarrow c$ .

- $(c_1 \text{ always}(\in) p')$  implies  $(p \text{ always}(\in) p')$ , where  $p \rightsquigarrow c_1$ .
- $(c_1 \text{ always}(=) p'.s)$  implies  $(p \text{ always}(=) p'.s)$ , where  $p \rightsquigarrow c_1$ .
- $(c.s_{11}...s_{1n} \text{ always}\bowtie(\subseteq) c.s_{21}...s_{2m})$  implies  $(p'.s_{11}...s_{1n} \text{ always}\bowtie(\subseteq) p'.s_{21}...s_{2m})$ , for all  $p'$  such that  $p' \rightsquigarrow c$ .
- $(p \text{ always}\bowtie(\in) p')$  implies  $(p \text{ always}(\in) p')$ .

#### 3.1.4 Digression: Regular Subset Inclusion Constraints

We end this section by briefly discussing regular subset inclusion-constraint macros, which cannot be converted to equivalent ULD constraints. We note that for the data models described in this chapter, regular inclusion constraints are not required. However, considering non-joined subset inclusion constraints highlights a possible extension to the ULD constraint language.

A general subset-inclusion macro takes the following form, where  $c_1$  and  $c_2$  are constructs,  $n \geq 0$ ,  $m \geq 0$ , and  $t$  is the name of the constraint.

constraint  $t$ : (  $c.s_{11}.s_{12}...s_{1n}$  always( $\subseteq$ )  $c.s_{21}.s_{22}...s_{2m}$  ).

For example, given collection constructs  $c_1$  and  $c_2$ , the macro ( $c_1$  always( $\subseteq$ )  $c_2$ ) states that for every instance  $i_1$  of  $c_1$  there is an instance  $i_2$  of  $c_2$  such that each member of  $i_1$  is also a member of  $i_2$ .

As noted, subset inclusion-constraint macros cannot be expressed in the ULD constraint language, but the previous constraint can be expressed in first-order logic as follows.

$$\forall x \exists y \forall z (c\text{-inst}(x, c_1) \wedge z \in x \Rightarrow c\text{-inst}(y, c_2) \wedge z \in y)$$

Notice that this expression is equivalent to the first-order logic expression:  $\forall x \exists y \neg \exists z (c\text{-inst}(x, c_1) \wedge z \in x \wedge \neg(c\text{-inst}(y, c_2) \wedge z \in y))$ , which can be read as “For every  $x$  there is a  $y$  that makes the statement ‘there is a  $z$  in  $x$  that is not in  $y$ ’ false.” To convert the first-order logic expression to an equivalent ULD constraint-language assertion requires the ability to iterate over (using negation) each  $z$  in a collection  $x$  for a particular  $y$ . However, we believe it is possible to support subset inclusion constraints by extending the ULD constraint language to include intensional predicates in the head of a constraint assertion, as shown in the following program. Note that such an extension does not require any modification to the Datalog $^-$  representation of ULD constraints presented in Chapter 2.

```
c-inst(Y, c1), subset(X, Y) ⇐ c-inst(X, c1).
subset(X, Y) ← c-inst(X, c1), c-inst(Y, c2), ¬notSubset(X, Y).
notSubset(X, Y) ← c-inst(X, c1), c-inst(Y, c2), Z ∈ X, ¬Z ∈ Y.
```

## 3.2 Traditional Database Data Models

This section gives ULD descriptions for traditional database data models. We present the relational data model in Section 3.2.1, with identical constructs as those in Example 1 of Chapter 2, but extended with constraints. We then consider a typical object-oriented

data model in Section 3.2.2. And finally, we describe the Entity-Relationship (E-R) data model [27] in Section 3.2.3, which combines some of the structures of the relational and object-oriented data models, but with slightly looser conformance constraints.

### 3.2.1 Relational Data Models

The relational data model was first introduced by Codd [33] and is the most widely used and implemented data model for databases [71]. The relational data-model constructs and constraints are given below (Data-Model 1). Note that we assume that a configuration stores information for a single relational database and that there is a special, atomic value *null* that is designated to represent the null value.

In the relational-model specification given below, the relation construct corresponds to a table definition, where each relation in a configuration has a unique name and a list of attribute definitions. A table serves as the extent for a relation. Each table must conform to exactly one relation, and a relation has exactly one table as an instance. Thus, tables and relations are one-to-one. Tables consist of tuples, and each tuple in a table contains a value for each attribute of the table's corresponding relation. We model a tuple as a list of values, where the  $i$ -th value is a value for the  $i$ -th attribute of the relation. The number of values in a tuple is precisely the number of attributes in its corresponding relation.

We consider the typical integrity constraints for the relation model, that is, primary and foreign keys. Each relation can have at most one primary key (**pKey**). A primary key designates one or more relation attributes as unique, that is, the combination of values that form the primary key must be unique for each tuple in a table. A foreign key (**fkey**) defines a referential integrity constraint in which each tuple in a table can reference a key-value in a target table. A foreign key constraint requires the designated values of the foreign-key attributes in a source relation to contain either corresponding values of the primary-key attributes in a target relation or *null* values.

As shown, **rel-1** through **rel-16** are constraints for the relational data model that can be expressed using macros. The first four constraints state that attribute lists, primary- and foreign-key attribute lists, and tuples are non-empty collections. Thus, each relation must have at least one attribute, and similarly, each tuple must have at least one value.

If a primary or foreign key is defined, it must contain at least one attribute. The next six constraints, **rel-5** through **rel-10**, define uniqueness restrictions. In particular, every relation in a configuration has a unique name, each attribute in an attribute list (including those for foreign and primary keys) has a unique name, each relation can have at most one primary key (that is, each primary key must have a unique **forRel** value), and each relation can have at most one foreign key for a particular set of attributes and target relation. Constraint **rel-11** requires attributes to be nested within attribute lists (that is, every attribute must be assigned to some relation). Constraints **rel-12** and **rel-13** require primary- and foreign-key attribute lists to always be associated with a primary and foreign key, respectively. Constraint **rel-14** requires tuples to always be contained within a table. The last two constraints (**rel-15** and **rel-16**) define the subset relationship between a primary (foreign) key's attribute list and the corresponding relation's attribute list.

The next three constraints, **rel-17** through **rel-19**, cannot be expressed using macros. Each of these constraints further defines the foreign-key construct. The first constraint (**rel-22**) requires each foreign key to reference (via the **toRel** selector) a relation with a primary key. The second constraint (**rel-23**) requires the number of attributes in the source relation designated by the foreign key to be the same as the number of attributes in the target relation's primary key. The last constraint (**rel-24**) requires the source attributes to have the same domain as their corresponding attributes in the target relation's primary key. We note that the attributes of the source relation designated by the foreign key are assumed to be ordered to match the order of the primary key attributes in the target relation.

The remaining constraints define conformance for the relational model. The first three conformance constraints define the relationship between tables, tuples, relations, and attributes. The last two constraints describe primary- and foreign-key restrictions on tuples. To model the foreign-key restriction, we use a constraint assertion that is not satisfiable; a positive *d-inst* formula appears in the body of the rule and a negative *d-inst* formula appears in the head of the rule, where both formulas have identical parameters. Note that the body of the assertion succeeds only if the foreign-key restriction is violated. Thus, whenever the body of the constraint holds the constraint assertion derives (or requires)

a contradiction. Therefore, although the assertion can never be satisfied, it accurately captures the semantics of the foreign-key restriction.

**Data-Model 1** (*The relational data-model constructs and constraints*)

```

construct relation      = {hasName->uldString, hasAtts->attList}
construct attList      = list of attribute
construct attribute    = {hasName->uldString, hasDomain->uldValuetype}
construct pKey         = {forRel->relation, keyAtts->pKeyAttList}
construct fKey         = {forRel->relation, toRel->relation, keyAtts->fKeyAttList}
construct pKeyAttList  = list of attribute
construct fKeyAttList  = list of attribute
construct table        = bag of tuple conf(domain=1,range=1):relation
construct tuple        = list of uldValue conf(domain=*,range=1):relation

```

```

constraint rel-1: ( attList isNonEmpty ).
constraint rel-2: ( pKeyAttList isNonEmpty ).
constraint rel-3: ( fKeyAttList isNonEmpty ).
constraint rel-4: ( tuple isNonEmpty ).

```

```

constraint rel-5: ( hasName keyOf relation ).
constraint rel-6: ( hasName keyOf attList ).
constraint rel-7: ( forRel keyOf pKey ).
constraint rel-8: ( forRel, toRel, keyAtts keyOf fKey ).
constraint rel-9: ( hasName keyOf pKeyAttList ).
constraint rel-10: ( hasName keyOf fKeyAttList ).

```

```

constraint rel-11: ( attList always(=) relation.hasAtts ).
constraint rel-12: ( attribute always(∈) attList ).
constraint rel-13: ( pKeyAttList always(=) pKey.keyAtts ).
constraint rel-14: ( fKeyAttList always(=) fKey.keyAtts ).
constraint rel-15: ( pKey.keyAtts always⋈(⊆) pKey.forRel.hasAtts ).
constraint rel-16: ( fKey.keyAtts always⋈(⊆) fKey.forRel.hasAtts ).

```

```

constraint rel-17: (A foreign key always refers to a target relation that has a primary key.)
c-inst(P, pKey), P.forRel:R ⇐ c-inst(F, fKey), F.toRel:R.

```

```

constraint rel-18: (Number of foreign key attributes must equal number of key attributes.)
|Ap|=N ⇐ c-inst(F, fKey), F.toRel:R, c-inst(P, pKey), P.forRel:R, P.keyAtts:Ap,

```

$F.\text{keyAtts}:A_f, |A_f|=N.$

**constraint rel-19:** (*Domain of foreign key attributes must equal domain of key attributes.*)

$A_p[I]=X_2, X_2.\text{hasDomain}:D \Leftarrow \text{c-inst}(F, \text{fKey}), F.\text{toRel}:R, \text{c-inst}(P, \text{pKey}), P.\text{forRel}:R,$   
 $P.\text{keyAtts}:A_p, F.\text{keyAtts}:A_f, A_f[I]=X_1, X_1.\text{hasDomain}:D.$

**constraint rel-conf1:** (*Every tuple in a table conforms to the corresponding relation*)

$\text{d-inst}(T, R) \Leftarrow \text{c-inst}(B, \text{table}), \text{c-inst}(R, \text{relation}), \text{d-inst}(B, R), T \in B.$

**constraint rel-conf2:** (*Tuple values conform to corresponding attribute domains*)

$\text{d-inst}(V, D) \Leftarrow \text{c-inst}(T, \text{tuple}), \text{c-inst}(R, \text{relation}), \text{d-inst}(T, R), R.\text{hasAtts}:A_R,$   
 $A_r[I]=A, A.\text{hasDomain}:D, T[I]=V.$

**constraint rel-conf3:** (*Tuples have the same number of values as relation attributes*)

$|T|=N \Leftarrow \text{c-inst}(T, \text{tuple}), \text{c-inst}(R, \text{relation}), \text{d-inst}(T, R), R.\text{hasAtts}:A_R, |A_r|=N.$

**constraint rel-conf4:** (*Primary key constraint*)

$T_1=T_2 \Leftarrow \text{c-inst}(T_1, \text{tuple}), \text{c-inst}(T_2, \text{tuple}), \text{d-inst}(T_1, R), \text{d-inst}(T_2, R),$   
 $\neg \text{differ}(T_1, T_2, R).$

$\text{differ}(T_1, T_2, R) \Leftarrow \text{c-inst}(T_1, \text{tuple}), \text{c-inst}(T_2, \text{tuple}), \text{d-inst}(T_1, R), \text{d-inst}(T_2, R),$   
 $\text{c-inst}(P, \text{pKey}), P.\text{forRel}:R, P.\text{keyAtts}:A_p, R.\text{hasAtts}:A_r, X_1 \in A_p, X_2 \in A_r,$   
 $X_1.\text{hasName}:N, X_2.\text{hasName}:N, A_r[I]=X_2, T_1[I]=V_1, T_2[I]=V_2, V_1 \neq V_2.$

**constraint rel-conf5:** (*Foreign key constraint*)

$\neg \text{d-inst}(T, R_s) \Leftarrow \text{c-inst}(F, \text{fKey}), F.\text{forRel}:R_s, \text{c-inst}(T, \text{tuple}), \text{d-inst}(T, R_s),$   
 $\neg \text{satisfyFKeyConstraint}(T, F).$

$\text{satisfyFKeyConstraint}(T, F) \Leftarrow \text{c-inst}(F, \text{fKey}), F.\text{forRel}:R_s, \text{c-inst}(T, \text{tuple}), \text{d-inst}(T, R_s),$   
 $\neg \text{hasNonNullFKeyAtts}(T, F).$

$\text{satisfyFKeyConstraint}(T, F) \Leftarrow \text{c-inst}(F, \text{fKey}), F.\text{forRel}:R_s, \text{c-inst}(T, \text{tuple}), \text{d-inst}(T, R_s),$   
 $\neg \text{hasNonMatchingTuple}(T, F).$

$\text{hasNonNullFKeyAtts}(T, F) \Leftarrow \text{c-inst}(F, \text{fKey}), F.\text{forRel}:R_s, \text{d-inst}(T, R_s), F.\text{keyAtts}:A_{L_f},$   
 $R_s.\text{hasAtts}:A_L, A_{L_f}[I]=A, A_f \in A_{L_f}, A.\text{hasName}:N, A_f.\text{hasName}:N, T[I] \neq \text{null}.$

$\text{hasNonMatchingTuple}(T, F) \Leftarrow \text{c-inst}(F, \text{fKey}), F.\text{forRel}:R_s, \text{c-inst}(T, \text{tuple}), \text{d-inst}(T, R_s),$   
 $F.\text{keyAtts}:A_{L_f}, R_s.\text{hasAtts}:A_L, A \in A_L, A_{L_f}[I]=A_f, A.\text{hasName}:N, A_f.\text{hasName}:N,$   
 $\neg \text{matchingAttributeVal}(T, I, F).$

$\text{matchingAttributeVal}(T_s, I, F) \Leftarrow \text{c-inst}(F, \text{fKey}), \text{c-inst}(P, \text{pKey}), F.\text{forRel}:R_s, F.\text{toRel}:R_t,$

$P.\text{forRel}:R_t, c\text{-inst}(T_s, \text{tuple}), d\text{-inst}(T, R_s), c\text{-inst}(T_t, \text{tuple}), d\text{-inst}(T, R_t),$   
 $R_s.\text{hasAtts}:AL_s, R_t.\text{hasAtts}:AL_t, AL_s[I_s]=A_s, AL_t[I_t]=A_t, F.\text{keyAtts}:AL_f, AL_f[I]=A_f,$   
 $A_f.\text{hasName}:N_f, A_s.\text{hasName}:N_f, P.\text{keyAtts}:AL_p, AL_p[I]=A_p, A_p.\text{hasName}:N_p,$   
 $A_t.\text{hasName}:N_p, T_s[I_s]=V, T_t[I_t]=V.$

### 3.2.2 Object-Oriented Data Models

Object-oriented database data models extend relational data models by adding structured data types and object identifiers (to distinguish among objects with equivalent values) [48]. We present a version of the  $O_2$  data model [48, 49] (Data-Model 2, below), which has played a significant role in shaping the ODMG Object Model [26]—an evolving standard for object-based data management.

We simplify the data model of  $O_2$  as follows. Class names denote implicit extents, that is, every class name can be used as a collection that contains its designated instances. Each class has a tuple-structured type that defines a set of attributes, each of which can be basic (that is, contain atomic types) or structured. A structured type is either another class or a set type, where a set type can be arbitrarily nested with structured types. (In the ODMG Object Model and in  $O_2$ , structured types include tuple structures in addition to class and set types). Class types imply an implicit subsumption hierarchy, which may be more complete than the user-specified class hierarchy (that is, there may be classes whose types are compatible, but are not designated as subclasses). We assume each object has an identifier and a value, which is either atomic or structured. We do not consider method signatures, although they can easily be added to the definition below.

In an  $O_2$  database, every object must belong to (that is, be an instance of) at least one class extent. An object is allowed to exist in more than one class extent as allowed by type subsumption. Finally, the  $O_2$  data model allows only single inheritance, although it can be extended to support multiple class inheritance [49].

We note that in both the ODMG Object Model and the  $O_2$  data model, attributes can contain arbitrarily nested tuple structures in addition to the structured types previously defined. However, when arbitrarily nested structures are considered, the corresponding restriction for checking type subsumption (*oodm-19* below) cannot be extended using the

ULD constraint language. That is, we cannot construct a stratified Datalog<sup>−</sup> program that accurately captures the class-superclass typing constraint for arbitrarily nested structures (as currently defined). In particular, an additional *compatibleAttribute* rule is needed to ensure that an attribute (variable  $A_2$  in *compatibleAttribute*) with a nested tuple type is present in the given tuple type (variable  $T_1$  in *compatibleAttribute*). Performing this check requires negation through the *incompatibleTypes* formula. Thus, *incompatibleTypes* would call a *compatibleAttribute* formula, and *compatibleAttribute* would call a negated *incompatibleTypes* formula. We note that, from a modeling perspective, little is gained by using nested tuple types, because classes, which can essentially represent named tuple types, can be arbitrarily nested.

**Data-Model 2** (*The  $O_2$  data-model constructs and constraints*)

```

construct class      = {hasName->uldString, hasType->tupleType, hasSuper->class}
construct tupleType  = set of attribute
construct attribute  = {hasName->uldString, hasType->nestedType}
construct nestedType = class | basicType | setType
construct basicType  = {hasType->uldValueType}
construct setType    = {ofType->nestedType}
construct object     = {hasId->uldString, hasValue->tupleValue}
                    conf(domain=*,range=+):class
construct tupleValue = set of property
construct property   = {hasName->uldString, hasValue->nestedValue}
construct nestedValue = object | basicValue | setValue
construct basicValue = {hasValue->uldValue}
construct setValue   = set of nestedValue

```

constraint oodm-1: ( tupleType isDistinct ).

constraint oodm-2: ( setType isDistinct ).

constraint oodm-3: ( tupleValue isDistinct ).

constraint oodm-4: ( setValue isDistinct ).

constraint oodm-5: ( hasName keyOf class ).

constraint oodm-6: ( hasName keyOf tupleType ).

constraint oodm-7: ( hasId keyOf object ).

constraint oodm-8: ( hasValue keyOf basicValue ).

constraint oodm-9: ( hasName keyOf tupleValue ).

constraint oodm-10: ( attribute always( $\in$ ) tupleType ).

constraint oodm-11: ( property always( $\in$ ) tupleValue ).

constraint oodm-12: (*Each tuple type must be connected to a class*)

$\text{c-inst}(C, \text{class}), C.\text{hasType}:T \Leftarrow \text{c-inst}(T, \text{tupleType}).$

constraint oodm-13: (*Each non-class nested type must be connected to a class*)

$\text{c-inst}(C, \text{class}), C.\text{hasType}:T \Leftarrow \text{c-inst}(T, \text{nestedType}), \neg \text{c-inst}(T, \text{class}),$   
 $\neg \text{hasParentType}(T).$

$\text{hasParentType}(T) \Leftarrow \text{c-inst}(T_s, \text{setType}), T_s.\text{ofType}:T.$

$\text{hasParentType}(T) \Leftarrow \text{c-inst}(T_a, \text{attribute}), T_a.\text{hasType}:T.$

constraint oodm-14: (*Each tuple value must be connected to an object*)

$\text{c-inst}(O, \text{object}), O.\text{hasValue}:T \Leftarrow \text{c-inst}(T, \text{tupleValue})$

constraint oodm-15: (*Each non-object nested value must be connected to an object*)

$\text{c-inst}(O, \text{object}), O.\text{hasValue}:V \Leftarrow \text{c-inst}(V, \text{nestedValue}), \neg \text{c-inst}(V, \text{object}),$   
 $\neg \text{hasParentValue}(V).$

$\text{hasParentValue}(V) \Leftarrow \text{c-inst}(V_s, \text{setValue}), V \in V_s.$

$\text{hasParentValue}(V) \Leftarrow \text{c-inst}(V_p, \text{property}), V_p.\text{hasValue}:V.$

constraint oodm-16: (*Subclass hierarchy is acyclic*)

$\neg C_1.\text{hasSuper}:C_2 \Leftarrow \text{c-inst}(C_1, \text{class}), \text{c-inst}(C_2, \text{class}), \text{subClassClosure}(C_2, C_1).$

$\text{subClassClosure}(C_1, C_1) \Leftarrow \text{c-inst}(C_1, \text{class}).$

$\text{subClassClosure}(C_1, C_3) \Leftarrow \text{c-inst}(C_1, \text{class}), \text{c-inst}(C_3, \text{class}), C_1.\text{hasSuper}:C_2,$   
 $\text{subClassClosure}(C_2, C_3).$

constraint oodm-17: (*Set types cannot be infinite*)

$T_1 \neq T_2 \Leftarrow \text{hasNestedSetType}(T_1, T_2).$

$\text{hasNestedSetType}(T_1, T_2) \Leftarrow \text{c-inst}(T_1, \text{setType}), \text{c-inst}(T_2, \text{setType}), T_1.\text{ofType}:T_2.$

$\text{hasNestedSetType}(T_1, T_3) \Leftarrow \text{c-inst}(T_1, \text{setType}), \text{c-inst}(T_3, \text{setType}), T_1.\text{ofType}:T_2,$   
 $\text{hasNestedSetType}(T_2, T_3).$

constraint oodm-18: (*Set values cannot be recursive*)

$V_1 \neq V_2 \Leftarrow \text{hasNestedSetValue}(V_1, V_2).$

$\text{hasNestedSetValue}(V_1, V_2) \leftarrow \text{c-inst}(V_1, \text{setValue}), \text{c-inst}(V_2, \text{setValue}), V_2 \in V_1.$   
 $\text{hasNestedSetValue}(V_1, V_3) \leftarrow \text{c-inst}(V_1, \text{setValue}), \text{c-inst}(V_3, \text{setValue}), V_2 \in V_1,$   
 $\text{hasNestedSetValue}(V_2, V_3).$

**constraint oodm-19:** (*Subclass type must be compatible with superclass type.*)

$\neg C_1.\text{hasSuper}:C_2 \Leftarrow \text{c-inst}(C_1, \text{class}), \text{c-inst}(C_2, \text{class}), C_1.\text{hasType}:T_1, C_2.\text{hasType}:T_2,$   
 $\text{incompatibleTypes}(T_1, T_2).$

$\text{incompatibleTypes}(T_1, T_2) \leftarrow \text{c-inst}(T_1, \text{tupleType}), \text{c-inst}(T_2, \text{tupleType}), A_2 \in T_2,$   
 $\neg \text{compatibleAttribute}(A_2, T_1).$

$\text{compatibleAttribute}(A_2, T_1) \leftarrow A_1 \in T_1, A_1.\text{hasName}:N, A_2.\text{hasName}:N, A_1.\text{hasType}:T'_1,$   
 $A_2.\text{hasType}:T'_2, \text{c-inst}(T'_1, \text{class}), \text{c-inst}(T'_2, \text{class}), \text{subClassClosure}(T'_1, T'_2).$

$\text{compatibleAttribute}(A_2, T_1) \leftarrow A_1 \in T_1, A_1.\text{hasName}:N, A_2.\text{hasName}:N, A_1.\text{hasType}:T'_1,$   
 $A_2.\text{hasType}:T'_2, \text{c-inst}(T'_1, \text{basicType}), \text{c-inst}(T'_2, \text{basicType}), T'_1.\text{hasType}:T,$   
 $T'_2.\text{hasType}:T.$

$\text{compatibleAttribute}(A_2, T_1) \leftarrow A_1 \in T_1, A_1.\text{hasName}:N, A_2.\text{hasName}:N, A_1.\text{hasType}:T'_1,$   
 $A_2.\text{hasType}:T'_2, \text{c-inst}(T'_1, \text{setType}), \text{c-inst}(T'_2, \text{setType}), \text{setSubType}(T'_1, T'_2).$

$\text{setSubType}(T_1, T_2) \leftarrow \text{c-inst}(T_1, \text{setType}), \text{c-inst}(T_2, \text{setType}), T_1.\text{ofType}:T'_1,$   
 $T_2.\text{ofType}:T'_2, \text{c-inst}(T'_1, \text{class}), \text{c-inst}(T'_2, \text{class}), \text{subClassClosure}(T'_1, T'_2).$

$\text{setSubType}(T_1, T_2) \leftarrow \text{c-inst}(T_1, \text{setType}), \text{c-inst}(T_2, \text{setType}), T_1.\text{ofType}:T'_1,$   
 $T_2.\text{ofType}:T'_2, \text{c-inst}(T'_1, \text{basicType}), \text{c-inst}(T'_2, \text{basicType}), T'_1.\text{hasType}:T,$   
 $T'_2.\text{hasType}:T.$

$\text{setSubType}(T_1, T_2) \leftarrow \text{c-inst}(T_1, \text{setType}), \text{c-inst}(T_2, \text{setType}), T_1.\text{ofType}:T'_1,$   
 $T_2.\text{ofType}:T'_2, \text{c-inst}(T'_1, \text{setType}), \text{c-inst}(T'_2, \text{setType}), \text{setSubType}(T'_1, T'_2).$

**constraint oodm-conf1:** (*Objects are instances of their corresponding superclasses*)

$\text{d-inst}(O, C_2) \Leftarrow \text{c-inst}(O, \text{object}), \text{d-inst}(O, C_1), C_1.\text{hasSuper}:C_2.$

**constraint oodm-conf2:** (*Objects must have compatible types to be class instances*)

$\neg \text{d-inst}(O, C) \Leftarrow \text{c-inst}(O, \text{object}), \text{c-inst}(C, \text{class}), O.\text{hasValue}:V, C.\text{hasType}:T,$   
 $\text{incompatibleValueForType}(V, T).$

$\text{incompatibleValueForType}(V, T) \leftarrow \text{c-inst}(V, \text{tupleValue}), \text{c-inst}(T, \text{tupleType}), A \in T,$   
 $\neg \text{hasPropertyForAttribute}(V, A).$

$\text{hasPropertyForAttribute}(V, A) \leftarrow \text{c-inst}(V, \text{tupleValue}), \text{c-inst}(A, \text{attribute}),$   
 $V.\text{hasName}:N, A.\text{hasName}:N, V.\text{hasValue}:V', A.\text{hasType}:T, \text{c-inst}(V', \text{object}),$   
 $\text{c-inst}(T, \text{class}), \text{d-inst}(V', T).$

$\text{hasPropertyForAttribute}(V, A) \leftarrow \text{c-inst}(V, \text{tupleValue}), \text{c-inst}(A, \text{attribute}),$

$V.\text{hasName:N}, A.\text{hasName:N}, V.\text{hasValue:V'}, A.\text{hasType:T},$   
 $c\text{-inst}(V', \text{basicValue}), c\text{-inst}(T, \text{basicType}), V'.\text{hasValue:X}, T.\text{hasType:T'},$   
 $d\text{-inst}(X, T').$   
 $\text{hasPropertyForAttribute}(V, A) \leftarrow c\text{-inst}(V, \text{tupleValue}), c\text{-inst}(A, \text{attribute}),$   
 $V.\text{hasName:N}, A.\text{hasName:N}, V.\text{hasValue:V'}, A.\text{hasType:T}, c\text{-inst}(V', \text{setValue}),$   
 $c\text{-inst}(T, \text{setType}), \neg \text{incompatValueForSetType}(V', T).$   
 $\text{incompatibleValueForSetType}(V, T) \leftarrow c\text{-inst}(V, \text{setValue}), c\text{-inst}(T, \text{setType}), X \in V,$   
 $T.\text{hasType:T'}, c\text{-inst}(X', \text{setValue}), \neg c\text{-inst}(T', \text{setType}).$   
 $\text{incompatibleValueForSetType}(V, T) \leftarrow c\text{-inst}(V, \text{setValue}), c\text{-inst}(T, \text{setType}), X \in V,$   
 $T.\text{hasType:T'}, c\text{-inst}(X, \text{basicValue}), \neg c\text{-inst}(T', \text{basicType}).$   
 $\text{incompatibleValueForSetType}(V, T) \leftarrow c\text{-inst}(V, \text{setValue}), c\text{-inst}(T, \text{setType}), X \in V,$   
 $T.\text{hasType:T'}, c\text{-inst}(X, \text{objectValue}), \neg c\text{-inst}(T', \text{class}).$   
 $\text{incompatibleValueForSetType}(V, T) \leftarrow c\text{-inst}(V, \text{setValue}), c\text{-inst}(T, \text{setType}), X \in V,$   
 $c\text{-inst}(X, \text{basicValue}), X.\text{hasValue:X'}, T.\text{hasType:T'}, c\text{-inst}(T', \text{basicType}),$   
 $\neg d\text{-inst}(X', T').$   
 $\text{incompatibleValueForSetType}(V, T) \leftarrow c\text{-inst}(V, \text{setValue}), c\text{-inst}(T, \text{setType}), X \in V,$   
 $c\text{-inst}(X, \text{object}), T.\text{hasType:T'}, c\text{-inst}(T', \text{class}), \neg d\text{-inst}(X, T').$   
 $\text{incompatibleValueForSetType}(V, T) \leftarrow c\text{-inst}(V, \text{setValue}), c\text{-inst}(T, \text{setType}), X \in V,$   
 $c\text{-inst}(X, \text{setValue}), T.\text{hasType:T'}, \text{incompatibleValueForSetType}(X, T').$

### 3.2.3 Entity-Relationship Data Models

The E-R data model [27] is primarily used as a conceptual modeling language for designing database schemas [71]. In particular, it can be used to define a conceptual schema that can then be mapped to relational table definitions (the logical schema) for storing data.

Information in the E-R data model is separated into entities and (explicit) relationships among entities. Both entities and relationships are grouped into types forming an E-R schema. An entity type consists of one or more attribute definitions. Attributes within an entity type can form keys, and each entity type must have exactly one primary key. A relationship type, which links two or more entity types, may also have attributes (although they are not required). Relationships do not have explicit keys. Instead, the primary key for a relationship is the combination of the primary keys of the entities being related. A relationship type, however, can be constrained using cardinality restrictions, which limit the participation of entities in the instances of the relationship type. A relationship type may also define role names for each entity type that participates in the relationship. The

role name represents the function an entity plays in the relationship and is useful for distinguishing the roles played in a recursive relationship type, that is, a relationship type that connects an entity type to itself.

An entity type can be designated as a *weak* entity type when its identity depends on another entity type. A weak entity type is specified using a combination of one or more of its attributes, called a *partial* key, and an *identifying* relationship type that connects the entity type to the entity type it depends on. Note that a partial key is required for a weak entity type. Thus, the primary key of a weak entity type consists of its partial key combined with its identifying entity type's primary key. A typical example of a weak entity is an employee dependent, where the dependent is covered by the employee's health insurance. Little information is recorded for each dependent; for example, a name and a birthdate may be the only attributes. The name of the dependent serves as its *partial* key and a relationship such as *dependent\_of* serves as its identifying relationship. A dependent entity is uniquely identified by specifying its name (the partial key) combined with an employee number (assuming the employee number is the primary key for an employee), where the dependent is connected to the employee through a *dependent\_of* relationship. Note that for each identifying entity-type instance, the instance's corresponding weak-entity instances are required to have distinct values for the partial key.

In the E-R data model, an entity is considered an instance of an entity type whenever it conforms to the specification of the type. Thus, the instance-of relationship in the E-R data model follows from the structure of entities, which is in contrast to object-oriented data models, where the user must explicitly specify the type of an object. It is also possible for an entity or relationship to exist without being classified by an available entity or relationship type, respectively. Most current applications of the E-R data model, however, assume schema is required, and only consider entities and relationships that are associated with types. This assumption is found in applications that use an E-R schema as a conceptual model for an underlying relational data-model representation.

We describe the E-R data model below (Data-Model 3). We assume that attributes are always single-valued (and not composite), that role names on relationship types are

optional, and that cardinality constraints restrict the participation of an entity in a relationship. In addition, we choose the general approach of not requiring entities and relationships to have corresponding types.

Note that a role is created for each actual role of a relationship type. Thus, for a relationship that contains two distinct roles for the same entity type, two instances of the role construct are created and stored for the role list. Therefore, role names and cardinality constraints can identify a particular role in a relationship type by giving a role identifier and a relationship-type identifier. We do not require roles to be unique across relationship types, that is, the same role can appear in multiple role lists. An entity type can also play the same role in multiple relationship types.

**Data-Model 3** (*The E-R data-model constructs and constraints*)

```

construct entityType = {hasName->uldString, hasAtts->attSet}
construct attSet     = set of attribute
construct attribute   = {hasName->uldString, hasDomain->uldValueType}
construct relType     = {hasName->uldString, hasRoles->roleList, hasAtts->attSet}
construct roleList    = list of role
construct role        = {hasType->entityType}
construct roleName    = {forType->relType, forRole->role, hasName->uldString}
construct pKey        = {forType->entityType, keyAtts->pKeyAttSet}
construct pKeyAttSet  = set of attribute
construct weakEntity  = {forType->entityType, idRel->relType,
                        partialKey->weakAttSet}
construct weakAttSet  = set of attribute
construct card        = {forType->relType, forRole->role, isReq->uldBool,
                        isMult->uldBool}
construct instance    = entity | rel
construct entity      = {hasProps->propSet} conf(domain=*,range=):entityType
construct propSet     = set of property
construct property    = {hasName->uldString, hasVal->uldValue}
construct rel         = {hasProps->propSet, associates->entityList}
                        conf(domain=*,range=):relType
construct entityList  = list of entity

constraint er-1: ( roleList isNonEmpty ).
constraint er-2: ( entityList isNonEmpty ).

```

constraint er-3: ( attSet isDistinct ).  
 constraint er-4: ( roleList isDistinct ).  
 constraint er-5: ( pKeyAttSet isDistinct ).  
 constraint er-6: ( weakAttSet isDistinct ).  
 constraint er-7: ( propSet isDistinct ).  
 constraint er-8: ( entityList isDistinct ).

constraint er-9: ( hasName keyOf entityType ).  
 constraint er-10: ( hasName keyOf attSet ).  
 constraint er-11: ( hasName keyOf relType ).  
 constraint er-12: ( forRel, forRole keyOf roleName ).  
 constraint er-13: ( forType keyOf pKey ).  
 constraint er-14: ( hasName keyOf pKeyAttSet ).  
 constraint er-15: ( forType keyOf weakEntity ).  
 constraint er-16: ( hasName keyOf weakAttSet ).  
 constraint er-17: ( forRel, forRole keyOf card ).  
 constraint er-18: ( hasName keyOf propSet ).

constraint er-19: ( attSet always(=) type.hasAtts ).  
 constraint er-20: ( roleList always(=) relType.hasRoles ).  
 constraint er-21: ( pKeyAttSet always(=) pKey.keyAtts ).  
 constraint er-22: ( weakAttSet always(=) weakEntity.partialKey ).  
 constraint er-23: ( propSet always(=) instance.hasProps ).  
 constraint er-24: ( entityList always(=) rel.associates ).  
 constraint er-25: ( attribute always( $\in$ ) attSet ).  
 constraint er-26: ( role always( $\in$ ) roleList ).  
 constraint er-27: ( property always( $\in$ ) propSet ).  
 constraint er-28: ( pKey.keyAtts always( $\bowtie$ ) ( $\subseteq$ ) pKey.forType.hasAtts ).  
 constraint er-29: ( weakEntity.partialKey always( $\bowtie$ ) ( $\subseteq$ ) weakEntity.forType.hasAtts ).  
 constraint er-30: ( roleName.forRole always( $\bowtie$ ) ( $\in$ ) roleName.forType.hasRoles ).  
 constraint er-31: ( card.forRole always( $\bowtie$ ) ( $\in$ ) card.forType.hasRoles ).

constraint er-32: (*Relationship types have unique sets of role names*)

$$R_{n1}=R_{n2} \Leftarrow \text{c-inst}(R_{n1}, \text{roleName}), \text{c-inst}(R_{n2}, \text{roleName}), R_{n1}.\text{forType}:R, R_{n2}.\text{forType}:R, \\ R_{n1}.\text{hasName}:N, R_{n2}.\text{hasName}:N.$$

constraint er-33: (*Role lists have at least two roles*)

$$R_1 \in R_l, R_2 \in R_l, R_1 \neq R_2 \Leftarrow \text{c-inst}(R_l, \text{roleList}).$$

constraint er-34: (*A role list contains unique role ids*)

$$R_1 \neq R_2 \Leftarrow \text{c-inst}(R_l, \text{roleList}), R_1 \in R_l, R_2 \in R_l.$$

constraint er-35: (*Every entity list has at least two elements*)

$$E_1 \in E_l, E_2 \in E_l, E_1 \neq E_2 \Leftarrow \text{c-inst}(E_l, \text{entityList}).$$

constraint er-36: (*Entity types have at least one attribute*)

$$A \in A_s \Leftarrow \text{c-inst}(E, \text{entityType}), E.\text{hasAtts}:A_s.$$

constraint er-37: (*Entities have at least one property*)

$$P \in P_l \Leftarrow \text{c-inst}(E, \text{entity}), E.\text{hasProps}:P_l.$$

constraint er-38: (*Weak entity types do not have primary keys*)

$$\neg P.\text{forType}:T \Leftarrow \text{c-inst}(T, \text{entityType}), \text{c-inst}(W, \text{weakEntity}), W.\text{forType}:T, \\ \text{c-inst}(P, \text{pKey}).$$

constraint er-39: (*Non-weak entity types always have primary keys*)

$$\text{c-inst}(P, \text{pKey}), P.\text{keyFor}:T \Leftarrow \text{c-inst}(T, \text{entityType}), \neg \text{isWeakEntityType}(T).$$

$$\text{isWeakEntityType}(T) \Leftarrow \text{c-inst}(T, \text{entityType}), \text{c-inst}(W, \text{weakEntity}), W.\text{forType}:T.$$

constraint er-40: (*Identifying relationships are binary and connect to weak entity types*)

$$|R_l|=2, L_1 \in R_l, L_1.\text{hasType}:E \Leftarrow \text{c-inst}(W, \text{weakEntity}), W.\text{forType}:E, W.\text{idRel}:R, \\ R.\text{hasRoles}:R_l.$$

constraint er-conf1: (*Entities have all attributes of their entity types*)

$$P \in P_l, P.\text{hasName}:N, P.\text{hasValue}:V, \text{d-inst}(V, D) \Leftarrow \text{c-inst}(T, \text{entityType}), \text{c-inst}(E, \text{entity}), \\ T.\text{hasAtts}:A_l, A \in A_l, A.\text{hasName}:N, A.\text{hasDomain}:D, E.\text{hasProps}:P_l.$$

constraint er-conf2: (*Relationships have all attributes of their relationship types*)

$$P \in P_l, P.\text{hasName}:N, P.\text{hasValue}:V, \text{d-inst}(V, D) \Leftarrow \text{c-inst}(T, \text{relType}), \text{d-inst}(R, T) \\ T.\text{hasAtts}:A_l, A \in A_l, A.\text{hasName}:N, A.\text{hasDomain}:D, R.\text{hasProps}:P_l.$$

constraint er-conf3: (*Relationships connect a valid entity for each relationship-type role*)

$$E_l[I]=E, \text{d-inst}(E, T_e) \Leftarrow \text{c-inst}(T_r, \text{relType}), \text{d-inst}(R, T), T_r.\text{hasRoles}:O_l, O_l[I]=O, \\ O.\text{hasType}:T_e, R.\text{associates}:E_l.$$

constraint er-conf4: (*Primary key constraint*)

$$E_1=E_2 \Leftarrow \text{c-inst}(E_1, \text{entity}), \text{c-inst}(E_2, \text{entity}), \text{d-inst}(E_1, T), \text{d-inst}(E_2, T),$$

$c\text{-inst}(pKey, P), P.\text{forType}:T, P.\text{keyAtts}:A_k, \neg\text{differOnAttribute}(E_1, E_2, A_k).$

$\text{differOnAttribute}(E_1, E_2, A_k) \leftarrow E_1.\text{hasProps}:P_{s1}, E_2.\text{hasProps}:P_{s2}, A \in A_k, P_1 \in P_{s1},$   
 $P_2 \in P_{s2}, A.\text{hasName}:N, P_1.\text{hasName}:N, P_2.\text{hasName}:N, P_1.\text{hasVal}:V_1, P_2.\text{hasVal}:V_2,$   
 $V_1 \neq V_2.$

**constraint er-conf5:** (*Relationship types have unique instances*)

$R_1 = R_2 \leftarrow c\text{-inst}(R_1, \text{rel}), c\text{-inst}(R_2, \text{rel}), d\text{-inst}(R_1, T), d\text{-inst}(R_2, T), T.\text{hasRoles}:R_l,$   
 $\neg\text{differOnRole}(R_1, R_2, R_l).$

$\neg\text{differOnRole}(R_1, R_2, R_l) \leftarrow c\text{-inst}(R_l, \text{roleList}), c\text{-inst}(R_1, \text{rel}), c\text{-inst}(R_2, \text{rel}), R_l[I] = O_l,$   
 $R_1.\text{associates}:E_{l1}, R_2.\text{associates}:E_{l2}, E_{l1}[I] = E_1,$   
 $E_{l2}[I] = E_2, E_1 \neq E_2.$

**constraint er-conf6:** (*Weak entities must have an identifying relationship*)

$d\text{-inst}(R, T_r), R.\text{associates}:E_l, E_l[I] = E \leftarrow c\text{-inst}(W, \text{weakEntity}), W.\text{forType}:T_e,$   
 $d\text{-inst}(E, T_e), W.\text{idRel}:T_r, T_r.\text{hasRoles}:R_l, R_l[I] = L, L.\text{hasType}:T_e.$

**constraint er-conf7:** (*Partial keys constraint*)

$E_1 = E_2 \leftarrow c\text{-inst}(W, \text{weakEntity}), W.\text{forType}:T_{e1}, W.\text{idRel}:T_r, d\text{-inst}(E_1, T_{e1}),$   
 $d\text{-inst}(E_2, T_{e1}), T_r.\text{hasRoles}:R_l, R_l[I_1] = L_1, L_1.\text{hasType}:T_{e1}, R_l[I_2] = L_2,$   
 $L_2.\text{hasType}:T_{e2}, T_{e1} \neq T_{e2}, d\text{-inst}(R_1, T_r), d\text{-inst}(R_2, T_r), R_1.\text{associates}:E_{l1},$   
 $E_{l1}[I_1] = E_1, R_2.\text{associates}:E_{l2}, E_{l2}[I_1] = E_2, E_{l1}[I_2] = E, E_{l2}[I_2] = E,$   
 $\neg\text{differOnPartialAttributes}(E_1, E_2, W).$

$\text{differOnPartialAttributes}(E_1, E_2, W) \leftarrow c\text{-inst}(W, \text{weakEntity}), c\text{-inst}(E_1, \text{entity}),$   
 $c\text{-inst}(E_2, \text{entity}), W.\text{partialKey}:A_s, A \in A_s, E_1.\text{hasProps}:P_{s1}, P_1 \in P_{s1},$   
 $E_2.\text{hasProps}:P_{s2}, P_2 \in P_{s2}, A.\text{hasName}:N, P_1.\text{hasName}:N, P_2.\text{hasName}:N,$   
 $P_1.\text{hasVal}:V_1, P_2.\text{hasVal}:V_2, V_1 \neq V_2.$

**constraint er-conf8:** (*Required cardinality constraint*)

$d\text{-inst}(R, T_r), R.\text{associates}:E_l, E_l[I] = E \leftarrow c\text{-inst}(T_r, \text{relType}), T_r.\text{hasRoles}:R_l, R_l[I] = L,$   
 $c\text{-inst}(C, \text{card}), C.\text{forRole}:L, C.\text{forType}:T_r, C.\text{isReq}:\text{true}, L.\text{hasType}:T_e,$   
 $c\text{-inst}(E, T_e).$

**constraint er-conf9:** (*Maximum cardinality constraint*)

$R_1 = R_2 \leftarrow c\text{-inst}(T_r, \text{relType}), T_r.\text{hasRoles}:R_l, R_l[I] = L, c\text{-inst}(C, \text{card}), C.\text{forRole}:L,$   
 $C.\text{forType}:T_r, C.\text{isMult}:\text{false}, L.\text{hasType}:T_e, c\text{-inst}(E, T_e), c\text{-inst}(R_1, T_r),$   
 $c\text{-inst}(R_2, T_r), R_1.\text{associates}:EL_1, R_2.\text{associates}:EL_2, EL_1[I] = E, EL_2[I] = E.$

## 3.3 Structured Data Models for the Web

In this section, we describe data models for XML and RDF, two standards for exchanging information over the Web. In Section 3.3.1, we extend the definition of the XML data model from Chapter 2 by including a full description of the Document Type Definition (DTD) schema language. In Section 3.3.2, we describe the full RDF data model with RDF Schema and briefly mention the flexibility of RDF for defining metamodels by allowing multiple levels of schema.

### 3.3.1 XML Data Models

The Extensible Markup Language (XML) [22] is a text format for representing and exchanging hierarchically structured information. Data in XML consists of elements, attributes, and character-string data. The data model of XML is hierarchical. An XML document is structured as a tree, where an internal node in the tree is an element, a leaf node in the tree is either an element or string content (called ‘PCDATA’), and the children of a node are ordered. An XML document must have exactly one root element. Every element has a tag name and can contain an optional set of attribute-value pairs.

We give a ULD description for the XML data model below, including a description of DTDs for defining schemas for XML documents. As shown, the majority of constructs introduced in the ULD description below are for representing the structure of DTDs—all but the last ten constructs describe the structure of DTDs.

With a DTD, it is possible to restrict the allowable elements and attributes in (conforming) XML documents. An element definition (`elementType`) in a DTD consists of a tag name and a structural restriction (`contentSpec`) on the element’s children (called a content model). The content model of an element is expressed in a DTD as a regular expression over element tag names. (Regular expressions are one reason DTDs are complex; the other reason is a result of the number of subtle restrictions possible for attributes.) In addition, each element definition in a DTD may also have a set of corresponding attribute definitions (`attDef`), which specify an attribute name and a restriction on the allowable attribute values.

An XML document specifies at most one DTD, which is an assertion that the document conforms to the structure defined by the DTD. Note that an XML document is not required to specify a DTD. By not providing a DTD, the document is considered un-typed, that is, the document does not have a specified schema (assuming no other schema language is used). Thus, the XML data model permits schema-less data.

An XML document is valid with respect to a DTD if the XML document: (1) only contains elements defined within the DTD, (2) each element only contains attributes defined in the DTD and the attributes contain valid values, (3) each required attribute defined in the DTD is present for each corresponding element in the document, and (4) the children nodes of an element are valid with respect to the corresponding element's content model defined in the DTD. Note that an element can be defined in a DTD such that it can contain any content, in which case the above rules are not considered for that particular element's children in a document. Thus, it is possible to mix typed and un-typed data in an XML document.

In a DTD, a content model is designated as either simple (**simpleSpec**), mixed (**mixedSpec**), or as a regular expression (**regExp**). There are three kinds of simple content models: an element is restricted to be empty, that is, it does not contain any children nodes; an element is restricted to contain only character data; or an element is restricted to contain any content including elements not specified in the DTD. A mixed content model allows an element to contain other elements mixed with character data, where the subelements are ordered according to their listing in the content-model specification. And finally, a regular-expression content model specifies the allowable organization of sub-elements for an element. Only elements defined in the DTD can occur as direct children of the element being specified by the regular expression (in contrast to a simple or mixed content model). A regular expression defines a grammar over elements and consists of operations applied to elements and sub-expressions. The allowable operations include concatenation expressed as a list (**listRE**), union expressed as a choice structure (**choiceRE**), and multiple forms of repetition including zero or more (**optManyRE**), zero or one (**optRE**), and one or more (**reqManyRE**).

Attributes can be constrained in a DTD to be required, implied, or fixed. A required

attribute must have a value for each corresponding element in the XML document. An implied attribute is optional. And a fixed attribute is required and must contain a specific value (given in the DTD) for each occurrence in the XML document. In addition, the allowable values of an attribute can be specified in a DTD. An attribute can either contain a character data value or be restricted to contain a value from an enumeration, a unique identifier, or a value from an attribute designated as an identifier (called an ‘IDREF’). We note that not all combinations of constraints and allowable values are possible, for example, an attribute cannot be defined in a DTD as both fixed and an identifier because each occurrence of an identifier (attribute) must be unique within an XML document.

We note that there are a number of features of DTDs that we do not consider below for simplicity, including attribute lists (in particular, IDREFS), name tokens, notations, entities, and default attributes. Many of these extensions are compiled away by an XML parser and do not significantly change the structural definition of DTDs. Although not included, these additional features have a straightforward definition using the ULD.

The definition below (Data-Model 4) defines both the constructs and constraints of the XML data model with DTDs. The constraints below, for example, for regular expression matching, define part of a validating parser for XML. That is, if all of the constraints are satisfied, the XML document is considered a model of the DTD. Constructing a complete validating parser from the definition below would require additional rules. Namely, we would need to infer the data-instance (*d-inst*) links for attributes and elements. However, these links are trivial to infer because they are based on element-tag and attribute names.

**Data-Model 4** (*The XML with DTD data-model constructs and constraints*)

```

construct dtd          = set of elemType
construct elemType     = {hasName->uldString, hasModel->contentSpec,
                           hasAtts->attDefList}
construct contentSpec  = simpleSpec | mixedSpec | regexp
construct simpleSpec   = {isPCData->uldBool, isEmpty->uldBool, isAny->uldBool}
construct mixedSpec    = list of elemType
construct regExp       = listRE | choiceRE | optManyRE | optRE | reqManyRE |
                           elemRE
construct listRE       = emptyRE | fullListRE

```

```

construct fullListRE    = {hasRE->regExp, hasTail->listRE}
construct choiceRE      = emptyRE | fullChoiceRE
construct fullChoiceRE  = {hasRE->regExp, hasTail->choiceRE}
construct emptyRE       = set of uldValue
construct optManyRE     = {hasRE->regExp}
construct optRE         = {hasRE->regExp}
construct reqManyRE     = {hasRE->regExp}
construct elemRE        = {hasType->elemType}
construct attDefList    = set of attDef
construct attDef        = enumAtt | cDataAtt | idAtt | idRefAtt
construct enumAtt       = reqEnumAtt | impEnumAtt
construct reqEnumAtt    = {hasName->uldString, hasEnum->enumSet}
construct impEnumAtt    = {hasName->uldString, hasEnum->enumSet}
construct enumSet       = set of cdata
construct cDataAtt      = reqCDataAtt | impCDataAtt | fixCDataAtt
construct reqCDataAtt   = {hasName->uldString}
construct impCDataAtt   = {hasName->uldString}
construct fixCDataAtt   = {hasName->uldString, hasFixVal->cdata}
construct idAtt         = reqIdAtt | impIdAtt
construct reqIdAtt      = {hasName->uldString}
construct impIdAtt      = {hasName->uldString}
construct idRefAtt      = reqIdRefAtt | impIdRefAtt | fixIdRefAtt
construct reqIdRefAtt   = {hasName->uldString}
construct impIdRefAtt   = {hasName->uldString}
construct fixIdRefAtt   = {hasName->uldString, hasFixVal->cdata}
construct doc           = {hasRoot->element} conf(domain=*,range=?):dtd
construct element       = {hasTag->uldString, hasAtts->attSet,
                           hasChildren->content} conf(domain=*,range=?):elemtype

construct attSet        = set of attribute
construct attribute     = {hasName->uldString, hasVal->cdata}
                           conf(domain=*,range=?):attDef
construct content       = emptyContent | fullContent
construct fullContent   = {hasNode->node, hasContent->content}
construct emptyContent  = set of uldValue
construct node          = pcdData | element
construct cdata         = atomic
construct pcdData       = atomic

constraint xml-1 ( dtd isNonEmpty ).

```

constraint xml-2 ( mixed isEmpty ).  
 constraint xml-3 ( enumSet isEmpty ).

constraint xml-4 ( hasName keyOf dtd ).  
 constraint xml-5 ( hasName keyOf elemType ).  
 constraint xml-6 ( hasName keyOf mixed ).  
 constraint xml-7 ( hasRE, hasListRE keyOf fullListRE ).  
 constraint xml-8 ( hasRE, hasChoiceRE keyOf fullChoiceRE ).  
 constraint xml-9 ( hasRE keyOf optManyRE ).  
 constraint xml-10 ( hasRE keyOf optRE ).  
 constraint xml-11 ( hasRE keyOf reqManyRE ).  
 constraint xml-12 ( hasType keyOf elemRE ).  
 constraint xml-13 ( hasName keyOf attDefList ).  
 constraint xml-14 ( hasName keyOf attSet ).  
 constraint xml-15 ( hasNode, hasContent keyOf fullContent ).

constraint xml-16 ( attDefList always(=) elemType.hasAtts ).  
 constraint xml-17 ( enumSet always(=) enumAtt.hasEnum ).  
 constraint xml-18 ( attSet always(=) element.hasAtts ).  
 constraint xml-19 ( emptyRE always(=) regExp.hasTail ).  
 constraint xml-20 ( emptyContent always(=) fullContent.hasContent ).  
 constraint xml-21 ( elemType always( $\in$ ) dtd ).  
 constraint xml-22 ( attDef always( $\in$ ) attDefList ).  
 constraint xml-23 ( attribute always( $\in$ ) attSet ).  
 constraint xml-24 ( pcdData always( $\in$ ) content ).

constraint xml-25 (*If an element is not a root, it is a child of another element*)  

$$\text{c-inst}(C, \text{fullContent}), C.\text{hasNode}:E \Leftarrow \text{c-inst}(D, \text{doc}), \text{c-inst}(E, \text{element})$$

$$\neg D.\text{hasRoot}:E.$$

constraint xml-26 (*If an element is not a child of another element, it is a root*)  

$$\text{c-inst}(D, \text{doc}), D.\text{hasRoot}:E \Leftarrow \text{c-inst}(E, \text{element}), \neg \text{childElement}(E).$$

$$\text{childElement}(E) \Leftarrow \text{c-inst}(E, \text{element}), \text{c-inst}(C, \text{fullContent}), C.\text{hasNode}:E.$$

constraint xml-27 (*Only one empty regular expression exists*)  

$$RE_1 = RE_2 \Leftarrow \text{c-inst}(RE_1, \text{emptyRE}), \text{c-inst}(RE_2, \text{emptyRE}).$$

constraint xml-28 (*The empty regular expression only occurs in the tail of a list or choice*)

$\neg \text{RE}_2.\text{hasRE}:\text{RE}_1 \Leftarrow \text{c-inst}(\text{RE}_1, \text{emptyRE}), \text{c-inst}(\text{RE}_2, \text{regExp}).$

constraint xml-29 (*The empty regular expression has no members*)

$\neg V \in \text{RE} \Leftarrow \text{c-inst}(\text{RE}, \text{emptyRE}), \text{c-inst}(V, \text{uldValue}).$

constraint xml-30 (*An empty content element has no members*)

$\neg V \in C \Leftarrow \text{c-inst}(C, \text{emptyContent}), \text{c-inst}(V, \text{uldValue}).$

constraint xml-31 (*Regular expression lists have at least two items*)

$\text{RE}_1.\text{hasTail}:\text{RE}_2, \text{c-inst}(\text{RE}_2, \text{fullListRE}) \Leftarrow \text{c-inst}(\text{RE}_1, \text{fullListRE}),$   
 $\neg \text{nestedListRE}(\text{RE}_1).$

$\text{nestedListRE}(\text{RE}_1) \Leftarrow \text{c-inst}(\text{RE}_1, \text{fullListRE}), \text{c-inst}(\text{RE}_2, \text{fullListRE}),$   
 $\text{RE}_2.\text{hasTail}:\text{RE}_1.$

constraint xml-32 (*Regular expression choices have at least two items*)

$\text{RE}_1.\text{hasTail}:\text{RE}_2, \text{c-inst}(\text{RE}_2, \text{fullChoiceRE}) \Leftarrow \text{c-inst}(\text{RE}_1, \text{fullChoiceRE}),$   
 $\neg \text{nestedChoiceRE}(\text{RE}_1).$

$\text{nestedChoiceRE}(\text{RE}_1) \Leftarrow \text{c-inst}(\text{RE}_1, \text{fullChoiceRE}), \text{c-inst}(\text{RE}_2, \text{fullChoiceRE}),$   
 $\text{RE}_2.\text{hasTail}:\text{RE}_1.$

constraint xml-33 (*Choice regular expressions contain unique items*)

$\text{RE} \neq \text{RE}' \Leftarrow \text{c-inst}(\text{RE}_1, \text{fullChoiceRE}), \text{RE}_1.\text{hasRE}:\text{RE}, \text{RE}_1.\text{hasTail}:\text{RE}_2,$   
 $\text{elementInChoiceRE}(\text{RE}', \text{RE}_2).$

$\text{elementInChoice}(\text{RE}', \text{RE}_2) \Leftarrow \text{RE}_2.\text{hasRE}:\text{RE}'.$

$\text{elementInChoice}(\text{RE}', \text{RE}_3) \Leftarrow \text{RE}_3.\text{hasTail}:\text{RE}_2, \text{elementInChoice}(\text{RE}', \text{RE}_2).$

constraint xml-34 (*Enumeration sets must have at least two items*)

$\text{CD}_1 \in \text{ES}, \text{CD}_2 \in \text{ES}, \text{CD}_1 \neq \text{CD}_2 \Leftarrow \text{c-inst}(\text{ES}, \text{enumSet}).$

constraint xml-35 (*Simple content is either pcdData, empty, or any*)

$\text{S.isAny}:\text{true} \Leftarrow \text{c-inst}(\text{S}, \text{simpleSpec}), \neg \text{S.isPCData}:\text{true}, \neg \text{S.isEmpty}:\text{true}.$

constraint xml-36 (*Simple content is either pcdData, empty, or any*)

$\text{S.isEmpty}:\text{true} \Leftarrow \text{c-inst}(\text{S}, \text{simpleSpec}), \neg \text{S.isPCData}:\text{true}, \neg \text{S.isAny}:\text{true}.$

constraint xml-37 (*Simple content is either pcdData, empty, or any*)

$\text{S.isPCData}:\text{true} \Leftarrow \text{c-inst}(\text{S}, \text{simpleSpec}), \neg \text{S.isEmpty}:\text{true}, \neg \text{S.isAny}:\text{true}.$

constraint xml-38 (*One DTD per configuration*)

$$D_1=D_2 \Leftarrow \text{c-inst}(D_1, \text{dtd}), \text{c-inst}(D_2, \text{dtd}).$$

constraint xml-39 (*One document per configuration*)

$$D_1=D_2 \Leftarrow \text{c-inst}(D_1, \text{doc}), \text{c-inst}(D_2, \text{doc}).$$

constraint xml-40 (*An element occurs only once in a document*)

$$E_1=E_2 \Leftarrow \text{c-inst}(E, \text{element}), E_1.\text{hasChildren}:C_1, E_2.\text{hasChildren}:C_2, E \in C_1, E \in C_2.$$

constraint xml-41 (*An element occurs only once in a document*)

$$\neg E \in C \Leftarrow \text{c-inst}(E, \text{element}), \text{c-inst}(D, \text{doc}), D.\text{hasRoot}:E, \text{c-inst}(E', \text{element}), \\ E'.\text{hasChildren}:C.$$

constraint xml-42 (*An element occurs only once in a document*)

$$\neg D.\text{hasRoot}:E \Leftarrow \text{c-inst}(E, \text{element}), \text{c-inst}(E', \text{element}), E'.\text{hasChildren}:C, E \in C, \\ \text{c-inst}(D, \text{doc}).$$

constraint xml-43 (*An attribute occurs only once in a document*)

$$E_1=E_2 \Leftarrow \text{c-inst}(A, \text{attribute}), \text{c-inst}(E_1, \text{element}), \text{c-inst}(E_2, \text{element}), E_1.\text{hasAtts}:AS_1, \\ E_2.\text{hasAtts}:AS_2, A \in AS_1, A \in AS_2.$$

constraint xml-44 (*Every configuration with an element has a document*)

$$\text{c-inst}(D, \text{doc}) \Leftarrow \text{c-inst}(E, \text{element}).$$

constraint xml-conf1 (*If an element has an element type, doc is an instance of the dtd*)

$$\text{c-inst}(T_D, \text{dtd}), \text{d-inst}(D, T_D) \Leftarrow \text{c-inst}(E, \text{element}), \text{d-inst}(E, T_E), \text{c-inst}(D, \text{doc}).$$

constraint xml-conf2 (*If an attribute has an element type, doc is an instance of the dtd*)

$$\text{c-inst}(T_D, \text{dtd}), \text{d-inst}(D, T_D) \Leftarrow \text{c-inst}(A, \text{attribute}), \text{d-inst}(A, T_A), \text{c-inst}(D, \text{doc}).$$

constraint xml-conf3 (*Attribute instances*)

$$N_A=N_T \Leftarrow \text{c-inst}(A, \text{attribute}), \text{d-inst}(A, T), A.\text{hasName}=N_A, T.\text{hasName}=N_T.$$

constraint xml-conf4 (*Enumerated attribute instances*)

$$V_A \in ES \Leftarrow \text{c-inst}(A, \text{attribute}), \text{d-inst}(A, T), \text{c-inst}(T, \text{enumAtt}), A.\text{hasVal}=V_A, \\ T.\text{hasEnum}=ES.$$

constraint xml-conf5 (*CData fixed attribute instances*)

$V_A = V_T \Leftarrow \text{c-inst}(A, \text{attribute}), \text{d-inst}(A, T), \text{c-inst}(T, \text{fixCDataAtt}), A.\text{hasVal} = V_A,$   
 $T.\text{hasFixVal} = V_T.$

**constraint xml-conf6** (*ID attribute instances*)

$A_1 = A_2 \Leftarrow \text{c-inst}(A_1, \text{attribute}), \text{c-inst}(A_2, \text{attribute}) \text{ d-inst}(A_1, T), \text{d-inst}(A_2, T),$   
 $\text{c-inst}(T, \text{idAtt}), A_1.\text{hasVal} = V, A_2.\text{hasVal} = V.$

**constraint xml-conf7** (*IDRef attribute instances*)

$\text{c-inst}(A_2, \text{attribute}), \text{d-inst}(A_2, \text{idAtt}), A_2.\text{hasVal} = V \Leftarrow \text{c-inst}(A_1, \text{attribute}), \text{d-inst}(A_1, T),$   
 $\text{c-inst}(T, \text{idRefAtt}), A_1.\text{hasVal} = V.$

**constraint xml-conf8** (*IDRef fixed attribute definitions*)

$V_A = V_T \Leftarrow \text{c-inst}(A, \text{attribute}), \text{d-inst}(A, T), \text{c-inst}(T, \text{fixIdRefAtt}), A.\text{hasVal} = V_A,$   
 $T.\text{hasFixVal} = V_T.$

**constraint xml-conf9** (*Element instances*)

$N_E = N_T \Leftarrow \text{c-inst}(E, \text{element}), \text{d-inst}(E, T), E.\text{hasTag} = N_E, T.\text{hasName} = N_T.$

**constraint xml-conf10** (*Elements must have all required element type attributes*)

$E.\text{hasAtts} = AS, A_E \in AS, \text{d-inst}(A_E, A_T) \Leftarrow \text{c-inst}(E, \text{element}), \text{d-inst}(E, T), T.\text{hasAtts} = AL,$   
 $A_T \in AL, \text{requiredAttDef}(A_T).$

$\text{requiredAttDef}(A_T) \Leftarrow \text{c-inst}(A_T, \text{reqEnumAtt}).$

$\text{requiredAttDef}(A_T) \Leftarrow \text{c-inst}(A_T, \text{reqCDataAtt}).$

$\text{requiredAttDef}(A_T) \Leftarrow \text{c-inst}(A_T, \text{reqIdAtt}).$

$\text{requiredAttDef}(A_T) \Leftarrow \text{c-inst}(A_T, \text{reqIdRefAtt}).$

**constraint xml-conf11** (*Elements may only have element type attributes*)

$A_T \in AL, A_T.\text{hasName} = N \Leftarrow \text{c-inst}(E, \text{element}), \text{d-inst}(E, T), T.\text{hasAtts} = AL, E.\text{hasAtts} = AS,$   
 $A_E \in AS, A_E.\text{hasName} = N.$

**constraint xml-conf12** (*Simple specification*)

$|C| = 0 \Leftarrow \text{c-inst}(E, \text{element}), \text{d-inst}(E, T), E.\text{hasChildren} = C, T.\text{hasModel} = M,$   
 $\text{c-inst}(M, \text{simpleSpec}), M.\text{isEmpty} = \text{true}$

**constraint xml-conf13** (*PCData specification*)

$V_1 = V_2, \text{c-inst}(V_1, \text{pcdata}) \Leftarrow \text{c-inst}(E, \text{element}), \text{d-inst}(E, T), E.\text{hasChildren} = C,$   
 $T.\text{hasModel} = M, \text{c-inst}(M, \text{simpleSpec}), M.\text{isPCData} = \text{true}, V_1 \in C, V_2 \in C.$

constraint xml-conf14 (*Mixed specification*)

$d\text{-inst}(V, T), T \in M \Leftarrow c\text{-inst}(E, \text{element}), d\text{-inst}(E, T), E.\text{hasChildren}:C, T.\text{hasModel}:M,$   
 $c\text{-inst}(M, \text{mixedSpec}), V \in C, \neg c\text{-inst}(V, \text{pdata}).$

constraint xml-conf15 (*Elements must have the same tagname as the element type*)

$E.\text{hasTag}:N, T.\text{hasName}:N \Leftarrow c\text{-inst}(E, \text{element}), d\text{-inst}(E, T).$

constraint xml-conf16 (*Regular expression specifications*)

$\neg d\text{-inst}(E, T) \Leftarrow c\text{-inst}(E, \text{element}), c\text{-inst}(T, \text{elemType}), E.\text{hasChildren}:C, T.\text{hasModel}:M,$   
 $\neg \text{regExpFor}(M, C).$

$\text{regExpFor}(M, C) \Leftarrow c\text{-inst}(M, \text{regExp}), c\text{-inst}(C, \text{content}), \text{regExpFor}(M, C, C'),$   
 $c\text{-inst}(C', \text{emptyContent}).$

$\text{regExpFor}(M, C, C) \Leftarrow c\text{-inst}(M, \text{emptyRE}), c\text{-inst}(C, \text{emptyContent}).$

$\text{regExpFor}(M, C, C') \Leftarrow c\text{-inst}(M, \text{elemRE}), c\text{-inst}(C, \text{fullContent}), M.\text{hasType}:T,$   
 $C.\text{hasNode}:E, d\text{-inst}(E, T), C.\text{hasContent}:C'.$

$\text{regExpFor}(M, C, C') \Leftarrow c\text{-inst}(M, \text{fullChoiceRE}), c\text{-inst}(C, \text{content}), M.\text{hasRE}:M',$   
 $\text{regExpFor}(M', C, C').$

$\text{regExpFor}(M, C, C') \Leftarrow c\text{-inst}(M, \text{fullChoiceRE}), c\text{-inst}(C, \text{content}), M.\text{hasTail}:M',$   
 $\text{regExpFor}(M', C).$

$\text{regExpFor}(M, C, C'') \Leftarrow c\text{-inst}(M, \text{fullListRE}), c\text{-inst}(C, \text{content}), M.\text{hasRE}:M',$   
 $\text{regExpFor}(M', C, C'), M.\text{hasTail}:M'', \text{regExpFor}(M, C', C'').$

$\text{regExpFor}(M, C, C) \Leftarrow c\text{-inst}(M, \text{optManyRE}), c\text{-inst}(C, \text{content}).$

$\text{regExpFor}(M, C, C') \Leftarrow c\text{-inst}(M, \text{optManyRE}), c\text{-inst}(C, \text{content}), M.\text{hasRE}:M',$   
 $\text{regExpFor}(M', C, C').$

$\text{regExpFor}(M, C, C'') \Leftarrow c\text{-inst}(M, \text{optManyRE}), c\text{-inst}(C, \text{content}), M.\text{hasRE}:M',$   
 $\text{regExpFor}(M', C, C'), \text{regExpFor}(M, C', C'').$

$\text{regExpFor}(M, C, C) \Leftarrow c\text{-inst}(M, \text{optRE}), c\text{-inst}(C, \text{content}).$

$\text{regExpFor}(M, C, C') \Leftarrow c\text{-inst}(M, \text{optRE}), c\text{-inst}(C, \text{content}), M.\text{hasRE}:M',$   
 $\text{regExpFor}(M', C, C').$

$\text{regExpFor}(M, C, C') \Leftarrow c\text{-inst}(M, \text{reqManyRE}), c\text{-inst}(C, \text{content}), M.\text{hasRE}:M',$   
 $\text{regExpFor}(M', C, C').$

$\text{regExpFor}(M, C, C'') \Leftarrow c\text{-inst}(M, \text{reqManyRE}), c\text{-inst}(C, \text{content}), M.\text{hasRE}:M',$   
 $\text{regExpFor}(M', C, C'), \text{regExpFor}(M, C', C'').$

constraint xml-conf17 (*Document instances*)

$$\text{d-inst}(E, T_E), T_E \in T_D \Leftarrow \text{c-inst}(D, \text{doc}), D.\text{hasRoot}:E, \text{d-inst}(D, T_D).$$

### 3.3.2 RDF Data Models

The Resource Description Framework (RDF) is a graph-based data model for exchanging information over the Web. Information in RDF can be represented in a number of formats, the most widely used is XML. The RDF data model serves as the foundation of the Semantic-Web initiative [11], which aims at creating an infrastructure to enable better search, integration, and discovery of Web-based content by adding semantic, machine processable information to existing Web resources. In the Semantic Web, RDF is seen as a language for exchanging semantic information among agents and other software systems.

The data model for RDF is based on semantic networks [53, 82] in which information is represented using directed graphs with labeled, binary edges and labeled nodes. In RDF, nodes are called resources and edges are called properties. A resource is denoted with a Uniform Resource Indicator (URI), which is a general form of a URL (therefore, a node in an RDF graph can denote a Web-page). Properties in RDF are also represented as URIs. A graph in RDF is represented by a set of triples, which take the form (*property*, *subject*, *object*), where the *subject* and *object* are resources, and the *property* connects the subject (as the domain of the property) to the object (as the range property).

Every assertion in an RDF document is made using triples. To create complex structures such as collections, RDF defines special properties and resources. For example, to represent list data structures, RDF uses the resource `rdf:Seq` (where `rdf` is a namespace and `Seq` is a fragment identifier) and properties `rdf:type` and `rdf:_1`, `rdf:_2`, etc. A resource  $r$  is defined as a list, called a *sequence* in RDF, by asserting the triple (`rdf:type`,  $r$ , `rdf:Seq`), which is read as “the type of  $r$  is a sequence” or equivalently as “ $r$  is an instance of a sequence.” Elements are added to a sequence using an assertion (`rdf:_1`,  $r$ ,  $r'$ ), which is read as “the first item of  $r$  is  $r'$ .” In addition, RDF has a set of special properties for making assertions about triples (called reification), allowing complex hyper-graph structures to be defined. Semantic network representations, such as RDF, provide highly extensible systems for modeling information at the cost of clarity and formalism

[82]. In particular, it is easy to build new complex structures from triples by defining special resources and properties. However, doing so often leads to confusion by those who must use the structures because certain properties and resources have a special status or are introduced as a convenience for modeling without being formally defined.

Schema languages for RDF are built using the flexibility offered by RDF. In particular, RDF Schema (RDFS), which is the base schema language for RDF, defines special resources and properties for classifying data. RDFS serves as a very simple description logic [7] that permits the definition of classes and their roles, with universal role restriction. In RDFS, the resource `rdfs:Class` denotes the concept of a class and the resource `rdfs:Property` denotes the concept of a property definition. A property definition includes a domain (through the special property `rdfs:domain`) and a range (through the special property `rdfs:range`) restriction. A particular class  $c$  is defined by asserting the triple `(rdf:type, c, rdfs:Class)`. Objects are created in a similar way by asserting a triple `(rdf:type, o, c)`, which is read as “the object  $o$  is an instance of class  $c$ .” A particular property definition (that is, an instance of `rdfs:Property`) can be used as a property in a triple. Finally, RDFS permits subclass (`rdfs:subClassOf`) and subproperty (`rdfs:subPropertyOf`) definitions. Recall that a property specifies a domain and a range class, where a triple that uses the property connects an instance of the domain class to an instance of the range class. A property  $p_1$  can be a subproperty of another property  $p_2$  if the domain class of  $p_1$  is a subclass of the domain class of  $p_2$ , and the range class of  $p_1$  is a subclass of the range class of  $p_2$ . Thus, a subproperty can restrict a property by requiring a more restrictive domain and range class.

We provide a description of the RDF with RDFS data model below (Data-Model 5). Instead of describing RDF as a standard semantic network representation, we use the ULD to promote RDFS class and property definitions to special constructs in the data model. This approach is closer to the approach of description logics, where schema and data reside at separate levels. Thus, in our representation below, classes and properties are explicitly defined, and can be independently accessed from the data they classify. Note, however, that our representation does not limit the expressibility of RDF with RDFS, because even though a class and a property are separate constructs, they are still considered resources,

and thus can still participate in additional assertions. The conformance definitions below were taken from both the RDF Model Theory [41] and the RDF Vocabulary Language [23]. We assume that an RDF processor computes the transitive closure of the subclass and subproperty relationships prior to the creation of the corresponding ULD configuration. Note however that the constraints themselves could be used to infer these additional relationships.

Finally, it is possible to define multiple levels of schema in RDF using the `rdf:type` property. In particular, the `rdf:type` property can be applied without bound and without violating conformance constraints. One use of creating multiple levels of schema is for defining metamodels (also called metaclasses or metaobjects). A metamodel defines an abstract schema for creating conceptual models. A number of systems support metamodeling, such as Telos [59, 43], UML [15, 42], CLOS [44], MOF [61], and Protege [60]. In RDFS, a metamodel is defined as a set of classes and properties with a specific intended semantics. Taking an example from Telos [59], we give a metamodel below using RDF(S) triples for defining conceptual models about document systems.

```
(rdf:type,    DocumentClass, rdfs:Class)
(rdf:type,    AgentClass,    rdfs:Class)
(rdf:type,    SimpleClass,    rdfs:Class)
(rdf:type,    source,         rdfs:Property)
(rdfs:domain, source,         DocumentClass)
(rdfs:range,  source,         AgentClass)
(rdfs:type,   content,        rdfs:Property)
(rdfs:domain, content,        DocumentClass)
(rdfs:range,  content,        SimpleClass)
```

In this example, `DocumentClass` represents a metaclass for classes concerning documents, for example, instances of `DocumentClass` could include the classes `JournalPaper` and `ConferencePaper`. An `AgentClass` is a metaclass grouping those classes that serve as the source of a `DocumentClass`. For example, a `Person` class might serve as the author of a `JournalPaper`, where `author` is a particular source property. The following triples define the `JournalPaper` and `Person` classes, along with a particular paper, using RDF(S) (the

example is again taken from Telos).

```
(rdf:type,      JournalPaper, DocumentClass)
(rdf:type,      Person,      AgentClass)
(rdf:type,      rdfs:Literal, SimpleClass)
(rdf:type,      author,      source)
(rdfs:domain,   author,      JournalPaper)
(rdfs:range,    author,      Person)
(rdf:type,      title,      content)
(rdfs:domain,   title,      JournalPaper)
(rdfs:range,    title,      rdfs:Literal)

(rdf:type,      Martian,     JournalPaper)
(rdf:type,      Stanley,     Person)
(rdf:type,      LaSalle,     Person)
(rdf:type,      Wang,        Person)
(author,        Martian,     Stanley)
(author,        Martian,     LaSalle)
(author,        Martian,     Wang)
(title,         Martian,     'The MARTIAN system')
```

Note that multiple levels of schema are used in this example. For example, Martian is an instance of JournalPaper, which is an instance of DocumentClass (which is an instance of rdfs:Class).

#### **Data-Model 5** (*The RDF(S) data-model constructs and constraints*)

```
construct resource    = rdfType | simpleRes | complexRes
construct rdfType     = class | property
construct class       = {hasURI->uldURI, hasLabel->uldString}
                      conf(domain=*,range=*) : class
construct property    = {hasURI->uldURI, hasDomain->class, hasRange->class}
                      conf(domain=*,range=*) : rdfType
construct subClass    = {hasSub->class, hasSuper->class}
construct subProp     = {hasSub->prop, hasSuper->prop}
construct simpleRes   = {hasURI->uldURI} conf(domain=*,range=*) : class
construct complexRes  = {hasURI->uldURI, hasVal->complexVal}
```

```

construct complexVal = triple | rdfBag | rdfSeq | rdfAlt
construct triple     = {hasPred->resource, hasSubj->resource, hasObj->objectVal}
construct objectVal  = resource | literal
construct rdfBag     = bag of object
construct rdfSeq     = list of object
construct rdfAlt     = set of object
construct literal    = atomic conf(domain:*,range:*) : class

```

```

constraint rdf-1: ( hasURI keyOf resource )
constraint rdf-2: ( hasLabel keyOf class )
constraint rdf-3: ( hasSuper, hasSub keyOf subClass )
constraint rdf-4: ( hasSuper, hasSub keyOf subProp )
constraint rdf-5: ( hasPred, hasSubj, hasObj keyOf triple )
constraint rdf-6: ( rdfBag isUnique )
constraint rdf-7: ( rdfSeq isUnique )
constraint rdf-8: ( rdfAlt isUnique )
constraint rdf-9: ( rdfAlt isEmpty )
constraint rdf-10: ( literal always(=) triple.hasObj )

```

```

constraint rdf-11: (Subclass relationship is transitive)
  c-inst(S3, subClass), S3.hasSub:C1, S3.hasSuper:C3  $\Leftarrow$  c-inst(S1, subClass),
  c-inst(S2, subClass), S1.hasSub:C1, S1.hasSuper:C2, S2.hasSub:C2, S2.hasSuper:C3.

```

```

constraint rdf-12: (Subproperty relationship is transitive)
  c-inst(S3, subProp), S3.hasSub:P1, S3.hasSuper:P3  $\Leftarrow$  c-inst(S1, subProp),
  c-inst(S2, subProp), S1.hasSub:P1, S1.hasSuper:P2, S2.hasSub:P2, S2.hasSuper:P3.

```

```

constraint rdf-13: (Property domains must be subclasses of super property domains)
  c-inst(Sc, subClass), Sc.hasSub:C1, Sc.hasSuper:C2  $\Leftarrow$  c-inst(Sp, subProp), Sp.hasSub:P1,
  Sp.hasSuper:P2, P1.hasDomain:C1, P2.hasDomain:C2, C1  $\neq$  C2.

```

```

constraint rdf-14: (Property ranges must be subclasses of super property ranges)
  c-inst(Sc, subClass), Sc.hasSub:C1, Sc.hasSuper:C2  $\Leftarrow$  c-inst(Sp, subProp), Sp.hasSub:P1,
  Sp.hasSuper:P2, P1.hasRange:C1, P2.hasRange:C2, C1  $\neq$  C2.

```

```

constraint rdf-conf1: (Classes with literal instances only have literal instances)
  c-inst(L2, literal)  $\Leftarrow$  c-inst(C, class), d-inst(L1, C), d-inst(L2, C), c-inst(L1, literal).

```

```

constraint rdf-conf2: (An instance of a class is also an instance of a super class)

```

$d\text{-inst}(X, C_2) \Leftarrow d\text{-inst}(X, C_1), c\text{-inst}(C_1, \text{class}), c\text{-inst}(S, \text{subClass}), S.\text{hasSub}:C_1,$   
 $S.\text{hasSuper}:C_2.$

**constraint rdf-conf3:** (*An instance of a property is also an instance of a super property*)

$d\text{-inst}(X, P_2) \Leftarrow d\text{-inst}(X, P_1), c\text{-inst}(P_1, \text{property}), c\text{-inst}(S, \text{subProp}), S.\text{hasSub}:P_1,$   
 $S.\text{hasSuper}:P_2.$

**constraint rdf-conf4:** (*Triples with constrained properties have conforming subjects and objects*)

$d\text{-inst}(S, C_d), d\text{-inst}(O, C_r) \Leftarrow c\text{-inst}(T, \text{triple}), T.\text{hasPred}:P, c\text{-inst}(P, \text{property}),$   
 $P.\text{hasDomain}:C_d, T.\text{hasSubj}:S, P.\text{hasRange}:C_r, T.\text{hasObj}:O.$

**constraint rdf-conf5:** (*Triples are entailed by super properties*)

$c\text{-inst}(T_2, \text{triple}), T_2.\text{hasPred}:P_2, T_2.\text{hasSubj}:S, T_2.\text{hasObj}:O \Leftarrow c\text{-inst}(T_1, \text{triple}),$   
 $T_1.\text{hasPred}:P_1, c\text{-inst}(S, \text{subProp}), S.\text{hasSub}:P_1, S.\text{hasSuper}:P_2, T_1.\text{hasSubj}:S,$   
 $T_1.\text{hasObj}:O.$

**constraint rdf-conf6:** (*Property instances have conforming domain and range values*)

$d\text{-inst}(P_{d1}, P_{d2}), d\text{-inst}(P_{r1}, P_{r2}) \Leftarrow d\text{-inst}(P_1, P_2), c\text{-inst}(P_2, \text{property}),$   
 $P_1.\text{hasDomain}:P_{d1}, P_1.\text{hasRange}:P_{r1}, P_2.\text{hasDomain}:P_{d2}, P_2.\text{hasRange}:P_{r2}.$

### 3.4 Additional Data Models

This section describes two additional data models for representing structured information. In Section 3.4.1, we describe the Topic Map data model, which like RDF, permits multiple levels of schema and is used to represent information on the Web. In Section 3.4.3, we describe a structured hypertext data model called VIKI, which explicitly permits (and encourages) un-typed and partially typed data. The term *emergent schema* is used with VIKI to describe its operational approach of allowing users to start with un-typed data, and as natural types “emerge,” types can be assigned to existing information. Both Topic Maps and VIKI were developed to make it easy for users to navigate unstructured information. A Topic Map provides structures and typing capabilities for navigating information stored externally, for example, in unstructured Web-pages or other documents. VIKI provides structures for organizing arbitrary text-based content that would reside externally, for example, in a Topic Map.

### 3.4.1 Topic-Map Data Models

Compared to RDF, the goal of the Topic-Map data model [13, 66] is to enable people (as opposed to software agents) to more easily locate and browse digital information. The Topic-Map data model contains structures similar to those found in traditional book indexes, that is, a possibly hierarchical list of subjects, each with one or more references to page numbers, see-also notes, and so on. Due to this similarity, Topic Maps are often referred to as “back-of-the-book” indexes [72] for digital resources. A Topic Map consists of a set of topics, occurrences, associations that connect topics, and scopes.

A topic is a concrete entity that serves to reify (or represent) exactly one abstract subject, for example, a subject may represent a concept in an ontology [13]. Topics contain zero or more base names and occurrences, and can be instances of other topics. A base name provides a convenient textual representation for a topic. An occurrence relates a resource with a topic, where a resource is either external (such as a Web-page), and denoted with a URI, or internal, and represented as a character string. An occurrence is an assertion stating that the topic is relevant to the resource. Finally, the topic instance-of relationship allows topics to be arranged into hierarchies. Both occurrences and associations (see below) can also be typed using topics, however, unlike topics they can have at most one topic type. Figure 3.1 shows a simple Topic Map consisting of a binary association between two topics *m1* and *a1*, where *m1* has two occurrences and *a1* has one occurrence. The remaining topics in Figure 3.1 are used as types.

A topic is said to play a role in an association. Each role in an association can be distinguished by an optional role name (similar to the E-R data model). A role name is assigned using a topic, that is, roles are filled by topics and are distinguished (named) by topics.

A scope is defined in a Topic Map as a special topic for labeling topic names, occurrences, and association roles. A scope establishes a context for applying the labeled item. For example, an *English* scope could be used to label the base names in the Topic Map that are in English. The scope could then be used whenever a user wishes to browse only




Figure 3.1: An example of a Topic Map represented graphically.

the English portion of the Topic Map. The default scope for a topic map is the *unconstrained scope*. Unless an explicit scope is given, the unconstrained scope is assumed. The *topic naming constraint* asserts that any two topics in the same scope with an identical base name implicitly refer to the same subject, and should be combined into a single topic (by a topic map processor).

The Topic Map data model is an ISO standard [13], and there are a number of versions of the model including an XML interchange format (XTM) [66]. In XTM, subjects and topics are decoupled to simplify the specification of a Topic Map. We describe a version of the XTM data model below (Data-Model 6). We assume each configuration represents a Topic Map. For convenience, we do not include the additional subject information provided by the full XTM syntax. We note that, while topics and subjects are still considered one-to-one in XTM, by making subjects explicit, it is possible to define a topic in multiple places within a Topic Map and then re-combine the definition into a single topic by an XTM processor. We assume the configuration is constructed after an XTM processor has parsed the Topic Map, and that the parser computes the transitive closure of the subtopic relationship.

Topic-Map data models place few restrictions on class-instance relationships, as reflected in the conformance definitions below. The main restrictions concern transitivity, namely, the class-instance relationship is not transitive, whereas the subtype-supertype relationship is. We include these assumptions below. There are some recent proposals for Topic-Map constraint languages including AsTMa [9] and Topic-Map templates [65] (which are themselves Topic Maps). In practice, if constraints are required in an application, they are currently expressed using conventions introduced by a top-level set of topics that all other topics are assumed to be instances of.

**Data-Model 6** (*The XTM data-model constructs and constraints*)

```

construct topic          = {hasId->uldString, hasBNames->baseNames,
                           hasOccurs->occurBag} conf(domain=*,range=*) :topic
construct subTopicOf     = {hasSubTopic->topic, hasSuperTopic->topic}
construct baseNames      = set of uldString
construct occurBag       = bag of occurrence
construct occurrence     = {hasValue->occurRes} conf(domain=*,range=?):topic
construct occurRes       = uldString | uldURI
construct association     = list of assocMember conf(domain=*,range=?):topic
construct assocMember    = {forAssoc->association, hasTopic->topic}
construct memberRole     = {forMember->assocMember, hasRole->topic}
construct bNameScope     = {forTopic->topic, forBName->uldString, hasScope->topic}
construct occurScope     = {forTopic->topic, forOccur->occurrence, hasScope->topic}
construct assocScope     = {forAssoc->association, forMember->assocMember,
                           hasScope->topic}

constraint xtm-1: ( hasId keyOf topic )
constraint xtm-2: ( hasSubTopic, hasSuperTopic keyOf subTopicOf )
constraint xtm-3: ( hasValue keyOf occurrence )
constraint xtm-4: ( forMember, hasRole keyOf memberRole )
constraint xtm-5: ( forTopic, forBName, hasScope keyOf bNameScope )
constraint xtm-6: ( forTopic, forOccur, hasScope keyOf occurScope )
constraint xtm-7: ( forAssoc, forMember, hasScope keyOf assocScope )

constraint xtm-8: ( baseNames always(=) topic.hasBNames )
constraint xtm-9: ( occurBag always(=) topic.hasOccurs )
constraint xtm-10: ( occurrence always(∈) occurBag )
constraint xtm-11: ( assocMember always⋈(∈) assocMember.forAssoc )

```

constraint xtm-12: ( bNameScope.forBName always $\bowtie$ ( $\in$ ) bNameScope.forTopic.hasBNames )

constraint xtm-13: ( occurScope.forOccur always $\bowtie$ ( $\in$ ) occurScope.forTopic.hasOccurs )

constraint xtm-14: ( assocScope.forMember always $\bowtie$ ( $\in$ ) assocScope.forAssoc )

constraint xtm-15: ( baseNames isUnique )

constraint xtm-16: ( association isNonEmpty )

constraint xtm-17: (*Topic naming constraint: topics have disjoint base names in a scope*)

$T_1 = T_2 \Leftarrow \text{c-inst}(S_1, \text{bNameScope}), \text{c-inst}(S_2, \text{bNameScope}), S_1.\text{hasScope}:T, S_2.\text{hasScope}:T,$   
 $S_1.\text{forTopic}:T_1, S_2.\text{forTopic}:T_2, S_1.\text{forBName}:B, S_2.\text{forBName}:B.$

constraint xtm-18: (*An association member is only in its corresponding association*)

$\neg M \in A_2 \Leftarrow \text{c-inst}(M, \text{assocMember}), M.\text{forAssoc}:A_1, \text{c-inst}(A_2, \text{association}), A_1 \neq A_2.$

constraint xtm-19: (*The subtopic relationship is transitive*)

$\text{c-inst}(S_3, \text{subTopicOf}), S_3.\text{hasSubTopic}:T_1, S_3.\text{hasSuperTopic}:T_3 \Leftarrow \text{c-inst}(S_1, \text{subTopicOf}),$   
 $\text{c-inst}(S_2, \text{subTopicOf}), S_1.\text{hasSubTopic}:T_1, S_1.\text{hasSuperTopic}:T_2, S_2.\text{hasSubTopic}:T_2,$   
 $S_2.\text{hasSuperTopic}:T_3.$

constraint xtm-conf1: (*Instances of a topic are also instances of the super topics*)

$\text{d-inst}(T, T_2) \Leftarrow \text{d-inst}(T, T_1), \text{c-inst}(S, \text{subTopicOf}), S.\text{hasSubTopic}:T_1, S.\text{hasSuperTopic}:T_2.$

### 3.4.2 The VIKI Data Model

The VIKI system [55, 76] is designed as a tool for creating and navigating hypertext documents. Broadly speaking, a hypertext document consists of text-based information organized for navigation. Most hypertext data models consist of nodes and links, where a node contains text and a link connects two or more nodes. For example, the World-Wide Web is considered a distributed hypertext. The data model in VIKI uses object, ordered collection, and composite primitive structures for representing hypertext documents. An object consists of zero or more slots, which are attribute-value pairs. All textual information is stored in object slots. Lists can be nested and can contain objects. And a composite, which is similar to a tuple, groups two or more objects or lists. In VIKI, a user constructs a hypertext by creating basic structures (objects, collections, and composites), by arranging the structures on the screen, and by defining visual characteristics for each

structure. For example, a user can create an object in VIKI, choose a shape and color for the object, and place it in a specific position on the screen.

Each basic structure in VIKI may have multiple associated types. For example, an object type can consist of a set of slot names, where a slot name is similar to an attribute name. Objects can contain any number of distinct slots, which consist of a name and a value for the slot. If for a particular object type, an object has a slot for every slot name defined in the type, the object can be assigned the object type. A collection type simply serves as a label, thus, a collection can be assigned any available collection type. A composite type consists of two or more component types (either object or collection types). A composite type, however, is strict. Composite types can only be used to type composite items with identical structure as that prescribed by the type. Namely, a composite structure conforms to a type if it has exactly one subelement for each component type. A type can define visual properties including a shape and a default color for instances. The bottom of Figure 3.2 shows a simple hypertext represented in the VIKI data model. The hypertext includes a composite containing an object and two collections, where each collection contains sub-objects. The top of Figure 3.2 gives a type description for the hypertext.

In VIKI, data structures are decoupled from their physical representation. For example, the same object can be represented multiple times on the screen and each representation can have a distinct set of physical properties. In addition, VIKI does not require the visual properties of basic structures to match the visual properties of their types.

VIKI allows users to explicitly type a particular structure defined in the hypertext. In addition, VIKI offers typing and structuring suggestions. For example, VIKI notifies the user when an item matches a type definition. VIKI also tries to recognize implied structure by suggesting collection and composite structures based on repeating patterns of structures, visual properties, and the proximity of items (via a spatial parser). For example, if objects are placed relatively close together and have similar characteristics (similar physical properties or attributes), then VIKI will suggest a collection for grouping the items.

We describe the VIKI data model constructs and constraints below. We do not include




Figure 3.2: An example of a hypertext and its type definition in VIKI.

the visual properties of structures and types (such as font, color, border width, etc.), but instead focus on the data-specific items. (We note that adding visual properties does not fundamentally change the conformance definitions or constructs defined below.) VIKI is an interesting data model to study because it is a mix between object-oriented data models (with objects and collections) and semi-structured data models (such as XML), that allows untyped, partially typed, and fully typed data.

**Data-Model 7** (*The VIKI data-model constructs and constraints*)

```

construct objType      = {hasName->uldString, hasSlotDefs->slotDefSet}
construct slotDefSet   = set of slotDef
construct slotDef      = {hasName->uldString}
construct collType     = {hasName->uldString}
construct compType     = {hasName->uldString, hasCompDef->compContDef}
construct compContDef  = list of baseType
construct baseType     = objType | collType
construct object       = set of objContent conf(domain=*,range=*) : objType
construct objContent   = uldString | slot

```

```

construct slot          = {hasName->uldString, hasText->uldString}
construct collection    = {hasName->uldString, hasContent->collContent}
                        conf(domain=*,range=*) : collType
construct collContent   = list of baseItem
construct baseItem      = object | collection
construct composite     = list of baseItem conf(domain=*,range=*) : compType

```

```

constraint viki-1: ( hasName keyOf objType )
constraint viki-2: ( hasName keyOf slotDef )
constraint viki-3: ( hasName keyOf collType )
constraint viki-4: ( hasName keyOf compType )
constraint viki-5: ( hasName, hasText keyOf slot )
constraint viki-6: ( hasName keyOf object )
constraint viki-7: ( hasName, hasContent keyOf collection )

```

```

constraint viki-8: ( slotDefSet always(=) objType.hasSlotDefs )
constraint viki-9: ( slotDef always(∈) slotDefSet )
constraint viki-10: ( compContDef always(=) compType.hasDef )
constraint viki-11: ( slot always(∈) object )
constraint viki-12: ( collContent always(=) collection.hasContent )

```

```

constraint viki-13: ( compContDef isUnique )
constraint viki-14: ( object isUnique )
constraint viki-15: ( collContent isUnique )
constraint viki-16: ( composite isUnique )
constraint viki-17: ( compContDef isNonEmpty )
constraint viki-18: ( composite isNonEmpty )

```

constraint viki-19: (*Collections are acyclic*)

$$C_1 \neq C_2 \Leftarrow \text{c-inst}(C_1, \text{collection}), \text{unnestCollectionContentOf}(C_1, C_2).$$

$$\text{unnestCollectionContentOf}(C_1, C_2) \Leftarrow \text{c-inst}(C_1, \text{collection}), C_1.\text{hasContent}:C_2, \\ \text{c-inst}(C_2, \text{collection}).$$

$$\text{unnestCollectionContentOf}(C_1, C_3) \Leftarrow \text{c-inst}(C_1, \text{collection}), C_1.\text{hasContent}:C_2, \\ \text{c-inst}(C_2, \text{collection}), \text{unnestCollectionContentOf}(C_2, C_3).$$

constraint viki-20: (*Composites have at least two items*)

$$B_1 \in C, B_2 \in C, B_1 \neq B_2 \Leftarrow \text{c-inst}(C, \text{composite}).$$

constraint viki-21: (*Composite types have at least two component types*)

$$T_1 \in D, T_2 \in D, T_1 \neq T_2 \Leftarrow \text{c-inst}(T, \text{compType}), T.\text{hasDef}:D.$$

constraint viki-conf1: (*Objects have all of their object-type attributes*)

$$S_o \in O, \text{c-inst}(S_o, \text{slot}), S_o.\text{hasName}:N \Leftarrow \text{d-inst}(O, T), \text{c-inst}(T, \text{objType}), \\ O.\text{hasSlotDefs}:D, S_t \in D, S_t.\text{hasName}:N.$$

constraint viki-conf2: (*Composites have the same number of items as their composite types*)

$$|C|=L \Leftarrow \text{d-inst}(C, T), \text{c-inst}(T, \text{compType}), T.\text{hasCompDef}:D, |D|=L.$$

constraint viki-conf3: (*Composite members conform to composite-type members*)

$$C[I]=B, \text{d-inst}(B, T') \Leftarrow \text{d-inst}(C, T), \text{c-inst}(T, \text{compType}), T.\text{hasCompDef}:D, D[I]=T'.$$

## 3.5 Properties of Conformance

This section introduces a number of properties for describing conformance, which we use to classify the data models described in this chapter. We consider both structural properties (for example, if conformance is required or optional) and process-oriented properties (for example, if conformance is explicitly created). By *process-oriented* we mean the process by which information is modeled or created within a data model. We conclude this chapter by discussing how the ULD compares to other data models in terms of its use of conformance.

### 3.5.1 Structural Conformance Properties

Table 3.1 compares the structural conformance properties for the relational, object-oriented, E-R, XML, RDF, Topic Map, and VIKI data models. We describe the conformance properties used in Table 3.1 below. Note that we use the phrase “an instance is typed by a schema” below to mean that a construct instance is an instance of (*d-inst*) another construct instance, the latter instance serving as a type in the schema.

**Schema Existence** For a data model, existence of a schema is *required* if every configuration of the data model requires a schema definition. Alternatively, existence of a schema is *optional* if a schema is not required (for all configurations of the data model).

Table 3.1: Structural conformance properties of data models.

| <i>Data Model</i> | <i>Schema<br/>Existence</i> | <i>Individual<br/>Participation</i> | <i>Extent<br/>Restriction</i> | <i>Individual<br/>Restriction</i> | <i>Schema<br/>Levels</i> |
|-------------------|-----------------------------|-------------------------------------|-------------------------------|-----------------------------------|--------------------------|
| Relational        | Required                    | Single                              | Complete                      | Complete                          | Single                   |
| O-O               | Required                    | Multiple                            | Complete                      | Minimal                           | Single                   |
| E-R               | Optional                    | Multiple                            | Mixed                         | Minimal                           | Single                   |
| XML               | Optional                    | Single                              | Mixed                         | Complete                          | Single                   |
| RDF(S)            | Optional                    | Multiple                            | Mixed                         | Minimal                           | Collapsed                |
| Topic Map         | Optional                    | Multiple/Single                     | Mixed                         | Not Applicable                    | Collapsed                |
| VIKI              | Optional                    | Multiple                            | Mixed                         | Minimal/NA/Complete               | Single                   |

**Individual Participation** A data model requires *single* participation of individuals if instances conform to at most one schema type (when a schema is present). Alternatively, a data model allows *multiple* participation of individuals if instances can conform to multiple schema types.

**Extent Restriction** By extent, we refer to the implied collection of instances of a construct within a configuration. A data model places a *complete* restriction on an extent if, when a schema exists, all instances in the extent are required to be typed by the schema. Alternatively, a data model allows a *mixed* extent if when a schema exists, instances in the extent can exist without being typed by the schema.

**Individual Restriction** A data model places a *complete* restriction on individual instances if every instance typed by a schema is precisely typed by the schema, that is, instances do not contain more or less information than defined by their type. A data model places a *minimal* restriction on individual instances if instances typed by a schema are allowed to contain additional information. We note that in some cases, data models do not place any restriction on the structure of typed, individual instances, in which case we say the individual restriction is *not applicable*.

**Schema Levels** A data model allows at most a *single* level of schema if schema types are not permitted to be instances of other types. Alternatively, a data model allows *multiple* levels of schema if schema types are permitted to be instances of other types. In addition, a data model is *collapsed* if it permits multiple levels of schema using a single construct (that is, it has a recursive conformance relationship to itself) and

the construct does not conform to any other constructs (for example, a topic in a Topic Map serves as both data and schema). A data model is *bounded* if it contains a schema construct (that is, a construct that some other construct can conform to) that does not conform to another construct. A bounded data model may or may not have multiple levels of schema. Bounded data models with multiple levels of schema, but without acyclic conformance definitions, have a fixed number of meta-schema levels. (Although possible, we do not consider data models without at least one conformance relationship.)

The schema existence property is reflected in the ULD by the range cardinality restriction on conformance relationships. If any conformance relationship in a data model has a required range constraint (that is, **range** is either 1 or +), schema is required for the data model. Alternatively, schema is considered optional for a data model if there are no required range constraints. For example, in the relational model, the table definition `table = bag of tuple conf(domain=1,range=1):relation` has a range cardinality constraint of 1, thus, the existence of schema is shown as required in Table 3.1. We apply the schema existence property to the data model as a whole, however, it is also possible to consider each conformance relationship separately.

Like schema existence, the individual participation property is also reflected in the ULD by the range cardinality constraint. In particular, if the value of the range constraint is \* or +, the data model permits multiple individual participation, and otherwise, permits at most single individual participation. Note that data models with more than one conformance definition (that is, more than one construct with a **conf** assertion) may have different individual participation properties for each such relationship.

The extent and individual restrictions in general are reflected in the ULD using constraint assertion rules. (As a special case, if the existence of schema is required and a particular conformance relationship has a required range constraint, we can infer a complete extent restriction for the schema construct.) Both the extent and individual restrictions are local to data model constructs: data models may have a different value of restriction for each conformance relationship. We note that the extent- and individual-restriction

properties are independent. In particular, a data model may require a complete extent restriction and allow a minimal individual restriction. For example, in object-oriented data models, every object must have an associated class (that is, for classes, the data model has a complete extent restriction), however associated objects may have additional attributes, not defined by the class.

The schema-levels property is global for a data model and is reflected in a data model's conformance relationships. In particular, whenever a construct  $C_1$  can conform to a construct  $C_2$ , and  $C_2$  can conform to a construct  $C_3$ , the data model permits multiple levels of schema (where  $C_1$ ,  $C_2$ , and  $C_3$  are not required to be distinct). When no such  $C_3$  exists, the data model permits only a single level of schema. All data models with a single level of schema are bounded, and collapsed data models are never bounded. Both the bounded and collapsed properties are reflected in the ULD description of a data model by directly analyzing constructs and their conformance definitions.

We note that, in some data models, structural conformance properties may be applied even if there is not an explicit conformance relationship established between constructs. For example, in RDF, the occurrence of a property in a triple places conformance-like restrictions on the subject and object of the triple even though there is not an explicit conformance definition between the constructs. (Note these restrictions are described using conformance assertions in the ULD.)

In Table 3.1, we only list multiple property values for a data model if it contains more than one conformance definition (including conformance-like assertions) and they have different conformance properties. In particular, a topic in a Topic Map can be the instance of multiple topics, but associations and occurrences are allowed to be the instance of at most one topic. In VIKI, an object type specifies a minimal individual restriction, a composite type requires a complete individual restriction, and the individual restriction for a collection type is not applicable.

Finally, we note that XML includes the special *any* content model, which enables a mixed extension restriction, that is, not all elements and attributes are required to be typed by a DTD. However, when a DTD is specified, every element and attribute not associated with an *any* content model must be completely constrained by the DTD.

### 3.5.2 Process-Oriented Conformance Properties

Table 3.2 compares the process-oriented conformance properties for the relational, object-oriented, E-R, XML, RDF, Topic Map, and VIKI data models. As shown, we consider how conformance is established, and in particular, how the *d-inst* relationship is created.

Conformance is established *implicitly* in a data model if it is generated as a side-effect of creating data or schema. For example, in the relational model, the creation of a relation automatically generates the relation’s corresponding table. Similarly, in many object-oriented models, an object is created using a class constructor, which establishes conformance from the object to the class (and to the appropriate superclasses). Note that data models that require conformance to be established *implicitly* are *schema-first*, because schema must exist prior to creating associated data.

We say conformance is *inferred* if it is computed. For example, in the original E-R data model [27], the type of an entity is computed based on its inherent structure. Often, inferred conformance definitions are computable upon request, for example, by initiating a parser (as in the case of XML or RDF).

Finally, conformance is established *explicitly* if the data model offers a mechanism for declaring conformance explicitly. For example, to create conformance in RDF(S), a triple must be asserted using the `rdf:type` property. We note that data models that allow conformance to be established explicitly must permit data to be defined prior to schema, that is, the data model is *schema-later*, and generally permit more flexible conformance definitions. For example, RDF, Topic Maps, and VIKI each require conformance to be established explicitly.

### 3.5.3 The ULD Conformance Properties

Here, we compare the ULD to other data models in terms of conformance. Even though the ULD is designed as a meta-data-model (that is, to represent other data models) and not as a general purpose data model, generally speaking, the ULD is another data model—it is a structured representation for data—that can represent constructs as well as schema and data.

Table 3.2: Process-oriented conformance properties of data models.

| <i>Data Model</i> | <i>How is conformance established?</i> |                                                               |
|-------------------|----------------------------------------|---------------------------------------------------------------|
| Relational        | Implicitly                             | When relations are defined                                    |
| O-O               | Implicitly                             | When an object is created                                     |
| E-R               | Inferred                               | An entity's type is computed from its structure               |
| XML               | Inferred                               | An XML processor checks conformance                           |
|                   | Explicitly                             | A document declares a document type                           |
| RDF(S)            | Inferred                               | Additional conformance computed from subclass and subproperty |
|                   | Explicitly                             | When the <code>rdf:type</code> property is used               |
| Topic Map         | Explicitly                             | Items are typed by stating their topic type                   |
| VIKI              | Inferred                               | Users can choose to accept the systems typing suggestions     |
|                   | Explicitly                             | Users can explicitly state the type of an item                |

For the structural conformance properties, the ULD requires the existence of schema (every data instance is an instance of a construct, and every construct is an instance of a construct type), permits multiple individual participation (via the union construct), requires complete extent restriction, and requires complete individual restriction. More importantly, the ULD is exactly a bounded data model that permits multiple levels of schema.

Note that none of the data models we consider are both bounded and permit multiple levels of schema. By being a bounded data model, the ULD can distinguish and enforce the construct type, construct, and data-instance levels and provide a clear definition for each level. Additionally, by allowing multiple levels of schema (according to the *d-inst* relationship), the ULD can represent data models that permit optional and multiple levels of schema.

Finally, in terms of process-oriented properties, conformance in the ULD is established implicitly (when data is created) using the *ct-inst* and *c-inst* relationship, and explicitly for the *d-inst* relationship. Note that, as defined in Chapter 2, the *d-inst* relationship is (explicitly) defined when an object is created using the *c-inst* relationship.

### 3.6 Summary

In this chapter, we used the ULD to define a number of data models, each with distinct constructs and conformance properties. We also extended the ULD constraint language

with constraint macros, which offer short-hand notations for expressing common data-model constraints. Each constraint macro has a corresponding expression in the ULD constraint language. (Thus, constraint macros can be executed using the ULD constraint language.) We used constraint macros and the ULD constraint language to specify the inherent constraints of each data model presented in this chapter. We also defined a classification scheme for data models based on conformance. Finally, we compared the conformance properties of the ULD to other data models.

The definitions of the data models presented in this chapter were taken from various sources. The relational and E-R data models are standard and are formalized in a number of textbooks. However, there are many slightly different object-oriented data-models, each with distinct features. We chose an early version of the O<sub>2</sub> data model, which offers many of the most common features found in object-oriented data models. The XML, RDF, and Topic-Map data models are being defined as standards. There has been some work aimed at formalizing certain features of XML and RDF. For instance, a partial axiomatization has been developed for RDF Schema [41], which we used directly to define conformance constraints. The VIKI data model was designed for use in a hypertext tool, and is described informally in research papers. In general, we believe the ULD can benefit a data-model engineer by offering a mechanism to formally define the constructs and constraints of a data model. By formalizing a data model, a data-model engineer can provide a clear, concise, and unambiguous specification of a data model.

Our goal in this chapter was to show that the ULD can be used to give complete descriptions of a wide range of data models. As future work, we believe that it would be useful to develop techniques to help data-model engineers verify that data-model constraints are consistent and correct (that is, correct with respect to either a formalization of a data model, or to an engineer's understanding of the data model). For example, it may be possible to generate sample schema and data instances from the ULD description of the data model, and ask the data-model engineer if the samples are valid. If a sample is not valid, we can infer that the constraints are not restrictive enough. Alternatively, it may be useful to generate samples that violate data-model constraints, and ask the data-model engineer if the samples are valid. If such a sample is considered valid, we can

infer that the ULD constraints are too restrictive. The challenge of this approach is to select the appropriate set of test cases (that is, samples) from the ULD definition of the data model to determine whether the data-model constraints are correct.

For the XML data model, we used this approach to verify that the constraints given are accurate. In particular, for our tests, we used Prolog to represent and execute the XML data-model constraints. We also (manually) defined a set of test documents for a specific DTD. The test documents were designed to highlight specific cases of the data-model constraints. For example, we defined both valid and invalid documents to test each regular-expression operator. In general, we believe the problem is similar to software testing, which may offer useful techniques to determine whether data-model constraints are correctly specified.

# Chapter 4

## Transformation

This chapter describes a language for data model, schema, and data transformation. The language extends the ULD query language defined in Chapter 2 and can express flexible mappings that access multiple levels of information including data model, schema, and instance data. We begin in Section 4.1 by describing the transformation language. In Section 4.2, we use the language to describe mappings from the E-R and XML data models to the relational data model, a schema mapping for the example schemas in Chapter 1, and a mapping that describes an RDF serialization in XML. Finally, in Section 4.3, we sketch approaches for reasoning with mappings expressed in the transformation language. In particular, we give examples that use constraints to check for well-formed transformations and to elaborate partial mappings.

### 4.1 The ULD Transformation Language

The ULD transformation language exploits stratified Datalog $\neg$  to express mappings between heterogeneous information sources. The ULD transformation language is similar to the Well-founded Object Language (WOL) [35] in that it is a declarative, logic-based language for expressing mappings. The WOL language, however, is designed specifically for a restricted object-oriented data model to express schema mappings, whereas our transformation language leverages the ULD to express a wide range of transformations.

A ULD *mapping rule* is expressed against source and target configurations. For example, the following rule converts E-R entity types to relations in the relational data model.

$$rel/c-inst(T, relation) \Leftarrow er/c-inst(T, entityType).$$

The rule is read “if  $T$  is an entity-type identifier in the  $er$  configuration, then  $T$  is a relation in the  $rel$  configuration,” where  $er$  is the source configuration and  $rel$  is the target configuration. We use the special symbol  $\Leftarrow$  to denote a mapping rule: the body of the rule is on the right side of the  $\Leftarrow$  symbol and the head of the rule is on the left side of the  $\Leftarrow$  symbol. The mapping above states only that entity-type identifiers in a source configuration are mapped to relation identifiers in a target configuration. Additional rules are needed to transform entity-type names to relation names, entity-type attributes to relation attributes, to convert primary keys, and so on.

A *transformation* consists of one or more mapping rules and an optional set of rules that define intensional predicates (used by the mapping rules). For present purposes, we assume there is exactly one source and one target configuration for a transformation. We also assume that a target configuration contains a description of its data model prior to the execution of a transformation. Target configurations may also contain additional schema and instance data.

A mapping rule takes the following form, where  $n \geq 1$ ,  $l \geq 0$  and  $m$  is the rule name.

mapping  $m$ :

$$h_1, h_2, \dots, h_n \Leftarrow b_1, b_2, \dots, b_l.$$

The head of a mapping rule consists of one or more positive ULD formulas  $h_1$  to  $h_n$  expressed over a target configuration. As in Datalog $\neg$ , the body of a rule consists of zero or more positive or negative extensional formulas  $b_1$  to  $b_l$  expressed against a source configuration. Optionally,  $b_1$  to  $b_l$  can also contain formulas expressed against a target configuration as well as intensional formulas defined in the transformation. The syntax for mapping rules is similar to ULD constraints, however, mapping rules have a different interpretation. Namely, a mapping rule is equivalent to a rule expressed using stratified Datalog $\neg$  extended to allow more than one formula in the head of the rule.

When a mapping rule is executed, each formula in the head of the rule is added to (or asserted in) a target configuration. In other words, a mapping rule specifies one or more

updates to target configurations. The output of a transformation is the updated target configuration that results from computing the fixpoint of the corresponding mapping and intensional rules.

In the previous example, a new relation is constructed in the target configuration for every entity type in the source configuration. As another example, the following mapping rule adds the name and an (empty) attribute list to every generated relation in *rel* (the target configuration).

$$\begin{aligned}
 &rel/T.hasName:N, rel/c-inst(N, uldString), rel/T.hasAtts:AS, \\
 &rel/c-inst(AS, attList) \Leftarrow \\
 &er/c-inst(T, entityType), er/T.hasName:N, er/T.hasAtts:AS.
 \end{aligned}$$

As shown, the head of this rule contains multiple formulas. Each formula in the head of the rule is asserted in the target configuration such that every variable is replaced with its corresponding variable assignment. (A new variable assignment is obtained in each evaluation of the body of the rule.) We allow multiple formulas in the head of a mapping rule primarily for convenience. However, there are certain mapping rules (which we describe in more detail below) that cannot be expressed in standard Datalog, that is, when only one formula is allowed in the head of a mapping rule.

The configuration name for a formula can be dropped in a mapping rule when default source and target configurations are specified. In particular, formulas in the body of a mapping rule without explicit configuration modifiers are matched against formulas in the default source configuration, and formulas in the head of the rule are inserted into the default target configuration. Thus, the previous example can be equivalently expressed as follows, where the default source configuration is *er* and the default target is *rel*.<sup>1</sup>

$$\begin{aligned}
 &T.hasName:N, c-inst(N, uldString), T.hasAtts:AS, c-inst(AS, attList) \Leftarrow \\
 &c-inst(T, entityType), T.hasName:N, T.hasAtts:AS.
 \end{aligned}$$

So far, we have assumed entity-type identifiers do not conflict with identifiers in the target configuration. That is, each entity-type identifier in the source configuration is

---

<sup>1</sup>Because we assume a single source and target configuration for the transformations presented in this chapter, we typically drop the configuration modifier in mapping rules.

added directly to the target configuration as an instance of a relation construct. If a source identifier is already used in the target configuration, the resulting configuration may not be well-formed. (For example, the identifier could become an instance of more than one non-union construct in the target.) The following rule guarantees that non-conflicting identifiers are added to the target. The rule is read “if T is an entity type in the source, generate a new identifier R in the target such that R is a relation.”

$$\text{c-inst}(\text{R}, \text{relation}) \Leftarrow \text{c-inst}(\text{T}, \text{entityType}).$$

This rule introduces an *implicit* mapping between identifiers T and generated identifiers R. If the rule contained a second formula in the head that also used R, the same variable assignment for R is used in both formulas (thus, the same value of R is used in both formulas to update the target). For example, the following rule uses two head formulas to connect the name of an entity type to the name of the corresponding relation. Note that an equivalent Datalog representation does not exist for this mapping rule. The mapping rule cannot be divided into an equivalent set of Datalog rules that contain exactly one formula in the head because R is assumed to be assigned the same value in both formulas.

$$\text{c-inst}(\text{R}, \text{relation}), \text{R.hasName:N} \Leftarrow \text{c-inst}(\text{T}, \text{entityType}), \text{T.hasName:N}.$$

We provide a mechanism to define *explicit* mappings between identifiers using the identifier-creation operator  $\mathcal{K}$ , which generates new identifiers using existing source-configuration identifiers (or using labels, which we describe below). The identifier-creation operator  $\mathcal{K}$  is a Skolem function [39]. An advantage of using a Skolem function to create identifiers is it provides a mechanism to retrieve the generated identifiers in subsequent rules. To illustrate, consider the following mapping rule that implicitly generates new identifiers.

$$\begin{aligned} \text{R.hasName:N}, \text{c-inst}(\text{N}, \text{uldString}), \text{R.hasAtts:AL}, \text{c-inst}(\text{AL}, \text{attList}) \Leftarrow \\ \text{c-inst}(\text{T}, \text{entityType}), \text{T.hasName:N}, \text{T.hasAtts:AS}. \end{aligned}$$

Once this rule is executed for a particular entity-type T, we can no longer add attributes of T to the corresponding relation R in subsequent rules. The connection between T and

R (and similarly for AS and AL) is lost. The following rule is an equivalent mapping that uses  $\mathcal{K}$  to define the explicit connection between T and R (and AS and AL). We describe the syntax of the  $\mathcal{K}$  operator below.

$$\begin{aligned} &R.\text{hasName:N}, \text{c-inst}(\text{N}, \text{uldString}), R.\text{hasAtts:AL}, \text{c-inst}(\text{AL}, \text{attList}) \Leftarrow \\ &\text{c-inst}(\text{T}, \text{entityType}), \text{T.hasName:N}, \text{T.hasAtts:AS}, \mathcal{K}^T(\text{R}), \mathcal{K}^{AS}(\text{AL}). \end{aligned}$$

In this example, the values (that is, variable assignments) of T are used to create identifiers that serve as variable assignments for R. (Note that T must be ground when  $\mathcal{K}$  is called.) Given the same value for T,  $\mathcal{K}^T(\text{R})$  always unifies R with the same identifier. Other mapping rules can now obtain the relation (or attribute list) identifier generated from the original entity type (or attribute set, respectively), as shown in the following mapping rule, which converts entity-type attributes to relation attributes.<sup>2</sup>

$$\begin{aligned} &A_r \in \text{AL}, \text{c-inst}(A_r, \text{attribute}), A_r.\text{hasName:N}, \text{c-inst}(\text{N}, \text{uldString}), \\ &A_r.\text{hasDomain:D} \Leftarrow \\ &\text{c-inst}(\text{T}, \text{entityType}), \text{T.hasAtts:AS}, A_e \in \text{AS}, A_e.\text{hasName:N}, \\ &A_e.\text{hasDomain:D}, \mathcal{K}^{AS}(\text{AL}), \mathcal{K}^{A_e}(A_r). \end{aligned}$$

In general, the identifier creation operator takes the form  $\mathcal{K}_L^{V_1, V_2, \dots, V_n}(\text{V})$  for  $n \geq 0$ , where  $V_1$  to  $V_n$  are (ground) variables representing source-configuration identifiers,  $L$  is an optional constant label, and  $V$  is a variable that is unified with the generated identifier. When  $V$  is bound, the call to  $\mathcal{K}$  returns true if  $V$  can be unified with the value of the created identifier, and false otherwise. The  $\mathcal{K}$  operator can equivalently be written as the Skolem function  $\mathcal{K}(V_1, \dots, V_n, L)$ . Thus, given values for  $V_1$  to  $V_n$  and  $L$ ,  $\mathcal{K}$  always returns the same value  $V$ .<sup>3</sup>

---

<sup>2</sup>Note that even with  $\mathcal{K}$ , there are still cases where multiple formulas in the head a mapping rule are needed. In particular, multiple formulas are needed if the body of the rule depends on information in the target configuration and the rule generates new information in the target. An example of this case is shown by the mapping rule `rdxfml-3` in Section 4.2.4.

<sup>3</sup>We implement the  $\mathcal{K}$  operator in Prolog (see Section 6.1) using the binary predicate `idgen`, where the first argument is a list of terms  $T$ , and the second argument is the generated identifier. The list of terms  $T$  is used to hold the ground, source-configuration identifiers ( $V_1$  to  $V_n$ ) and the constant label ( $L$ ). For example, the expression  $\mathcal{K}_{attList}^{AS}(\text{AL})$  is represented in Prolog as `idgen([AS, 'attList'], AL)`.

The expression  $A_r \in AL$  in the head of the previous example inserts  $A_r$  into the list  $AL$ . In general, an expression  $X \in C$  in the head of a mapping rule serves to insert the value  $X$  into the collection  $C$  for each unique evaluation of the body of the rule. For lists, we assume the expression  $X \in C$  appends  $X$  to the end of the list. It is also possible to state the index of  $X$  in the list  $C$  explicitly using the expression  $C[I] = X$ , which also adds  $X$  to  $C$ .

In addition to the  $\mathcal{K}$  operator, we also permit other special-purpose functions. For example, mathematical operations such as addition as well as concatenation of strings (denoted using the ‘.’ operator) are used in the examples of Section 4.2. We also permit user-defined functions, which are often used to resolve differences in data. For example, a user-defined function can be used to convert full names in last-name-first format to full names in first-name-first format. As in Datalog, we require function parameters to be ground when a function is called, that is, every parameter is either a constant or a variable unified with a constant through a variable binding.

Finally, we note that it is possible to create non-terminating mapping rules using the ULD transformation language. For instance, consider the following rule (which does not terminate), where  $r$  is the name of the target configuration and  $c$  is a construct.

$$r/c\text{-inst}(Y, c) \Leftarrow r/c\text{-inst}(X, c), \mathcal{K}^X(Y).$$

This rule does not have a fixpoint: Each time the body is evaluated, a new identifier is generated, and the formula inserted into the target causes the body of the rule to have a new evaluation. The rule results in infinite recursion. (Note that we can restrict the domain of  $\mathcal{K}$  to source-configuration variables, in which case, the rule above would not be well-formed.)

## 4.2 Transformation Examples

This section defines a set of common transformations using the ULD transformation language. We start by describing an E-R to relational transformation. We divide the transformation into a set of schema and data mappings. By applying both transformations, the

schema and instance data in an E-R configuration is converted to a relational schema and a corresponding database in the target configuration. Alternatively, the schema transformation can be applied in isolation, which is what is typically done in conceptual modeling approaches. We then give two standard XML-to-relational transformations. The first transformation converts XML instance data directly to a relational database, without the use of an XML DTD. The second transformation requires the XML source to have a DTD, which is converted to a relational schema, and the XML instance data is converted to a corresponding database for the schema. Next, we describe a schema transformation from the XML DTD of Figure 1.2 to the RDF Schema of Figure 1.3. The transformation can be used to convert XML documents that conform to the DTD in Figure 1.2 to RDF instances that conform to the RDF Schema in Figure 1.3. Finally, we describe a transformation from RDF to XML that defines an XML interchange format for RDF similar to the standard RDF XML syntax [47]. The advantage of such a transformation is that the mapping between the data model and its serialization syntax (that is, the syntax used to interchange schema and instance data in the data model) is made explicit.

We note that none of the transformations described here can be fully expressed in traditional meta-data-model frameworks. In particular, the transformations use a mix of data-model, schema, and instance data to introduce conventions for representing a source configuration in a target data model.

#### 4.2.1 E-R to Relational Transformations

The E-R to relational schema mapping below is similar to the mapping given by Elmasri and Navathe [38]. Each entity type is converted to a relation (**er-rel-1**) and the attributes of entity types are mapped to corresponding relation attributes (**er-rel-2**). Similarly, each relationship type is mapped to a relation (**er-rel-4**) and relationship-type attributes are mapped to relation attributes (**er-rel-5**).

We simplify the transformation defined by Elmasri and Navathe as follows. Instead of using the primary key of an entity type as the primary key for the corresponding relation, we introduce a new relation attribute called “id” (**er-rel-3**). The id-attribute serves as a key for the relation. Each relation that is created from an E-R relationship type contains one

attribute for each role of the relationship type (er-rel-6). A foreign key is also created for each role attribute to the corresponding id-attribute. (We note that it is possible to define a mapping using the ULD transformation language to convert entity-type primary keys to relation primary keys, but for exposition, we choose the simpler id-attribute approach below.)

**Transformation 1** (*E-R-to relational schema transformation*)

mapping er-rel-1: (*Entity types become relations*)

$$\begin{aligned} & \text{c-inst}(\text{R}, \text{relation}), \text{R.hasName:N}, \text{c-inst}(\text{N}, \text{uldString}) \Leftarrow \\ & \text{c-inst}(\text{T}, \text{entityType}), \text{T.hasName:N}, \mathcal{K}_{\text{relation}}^T(\text{R}). \end{aligned}$$

mapping er-rel-2: (*Entity-type attributes become relation attributes*)

$$\begin{aligned} & \text{R.hasAtts:AL}, \text{c-inst}(\text{AL}, \text{attList}), \text{c-inst}(\text{A}_r, \text{attribute}), \text{A}_r \in \text{AL}, \text{A}_r.\text{hasName:N}, \\ & \text{c-inst}(\text{N}, \text{uldString}) \text{ A}_r.\text{hasDomain:D} \Leftarrow \\ & \text{c-inst}(\text{T}, \text{entityType}), \text{T.hasAtts:AS}, \text{A}_e \in \text{AS}, \text{A}_e.\text{hasName:N}, \text{A}_e.\text{hasDomain:D}, \\ & \mathcal{K}_{\text{relation}}^T(\text{R}), \mathcal{K}_{\text{attList}}^{\text{AS}}(\text{AL}), \mathcal{K}_{\text{attribute}}^{\text{A}_e}(\text{A}_r). \end{aligned}$$

mapping er-rel-3: (*Create a key for each relation*)

$$\begin{aligned} & \text{c-inst}(\text{A}, \text{attribute}), \text{A.hasName:'id'}, \text{c-inst}(\text{'id'}, \text{uldString}), \text{A.hasDomain:'uldString'}, \\ & \text{A} \in \text{AL}, \text{c-inst}(\text{P}, \text{pKey}), \text{P.forType:R}, \text{P.keyAtts:PL}, \text{c-inst}(\text{PL}, \text{pKeyAttList}), \text{A} \in \text{PL} \Leftarrow \\ & \text{c-inst}(\text{T}, \text{entityType}), \text{T.hasAtts:AS}, \mathcal{K}_{\text{relation}}^T(\text{R}), \mathcal{K}_{\text{attList}}^{\text{AS}}(\text{AL}), \mathcal{K}_{\text{id}}^T(\text{A}), \mathcal{K}_{\text{pkey}}^T(\text{P}), \\ & \mathcal{K}_{\text{pkeyattlist}}^T(\text{PL}). \end{aligned}$$

mapping er-rel-4: (*Relationship types become relations.*)

$$\begin{aligned} & \text{c-inst}(\text{R}, \text{relation}), \text{R.hasName:N}, \text{c-inst}(\text{N}, \text{uldString}) \Leftarrow \\ & \text{c-inst}(\text{T}, \text{relType}), \text{T.hasName:N}, \mathcal{K}_{\text{relation}}^T(\text{R}). \end{aligned}$$

mapping er-rel-5: (*Relationship attributes become relation attributes.*)

$$\begin{aligned} & \text{R.hasAtts:AL}, \text{c-inst}(\text{AL}, \text{attList}), \text{c-inst}(\text{A}_r, \text{attribute}), \text{A}_r \in \text{AL}, \text{A}_r.\text{hasName:N}, \\ & \text{c-inst}(\text{N}, \text{uldString}) \text{ A}_r.\text{hasDomain:D} \Leftarrow \\ & \text{c-inst}(\text{T}, \text{relType}), \text{T.hasAtts:AS}, \text{A}_e \in \text{AS}, \text{A}_e.\text{hasName:N}, \text{A}_e.\text{hasDomain:D}, \\ & \mathcal{K}_{\text{relation}}^T(\text{R}), \mathcal{K}_{\text{attList}}^{\text{AS}}(\text{AL}), \mathcal{K}_{\text{attribute}}^{\text{A}_e}(\text{A}_r). \end{aligned}$$

mapping er-rel-6: (*Relationship roles become foreign key attributes.*)

$$\begin{aligned} & \text{c-inst}(\text{A}, \text{attribute}), \text{A} \in \text{AL}, \text{A.hasName:N}, \text{c-inst}(\text{N}, \text{uldString}), \text{A.hasDomain:'uldString'}, \\ & \text{c-inst}(\text{F}, \text{fKey}), \text{fKey.forRel:R}_r, \text{fKey.toRel:R}_e, \text{fKey.keyAtts:FL}, \text{c-inst}(\text{FL}, \text{fKeyAttList}), \\ & \text{A} \in \text{FL} \Leftarrow \end{aligned}$$

$$\begin{aligned} & \text{c-inst}(\text{T}_r, \text{relType}), \text{T}_r.\text{hasAtts:AS}, \text{T}_r.\text{hasRoles:RL}, \text{RL}[\text{I}]=\text{R}, \text{R}.\text{hasType:T}_e, \\ & \text{N}=\text{'role'} \cdot \text{I}, \mathcal{K}_{\text{attribute}}^{R,I}(\text{A}), \mathcal{K}_{\text{fkey}}^{R,I}(\text{F}), \mathcal{K}_{\text{fkeyattlist}}^{R,I}(\text{FL}), \mathcal{K}_{\text{attList}}^{\text{AS}}(\text{AL}), \mathcal{K}_{\text{relation}}^{\text{T}_r}(\text{R}_r), \\ & \mathcal{K}_{\text{relation}}^{\text{T}_e}(\text{R}_e). \end{aligned}$$

The E-R to relational data mapping given below stores each entity and relationship as a tuple in one or more relations. For example, every entity that has an entity type is represented as a tuple in the entity type's corresponding relation. An entity or relationship with more than one type is transformed into multiple tuples, one for each corresponding type.

The data transformation converts only those entities and relationships with a corresponding entity and relationship type. Thus, the mapping selects only the E-R items that are covered by schema. We could additionally represent those entities and relationships that do not have corresponding schema by defining two default relations, one for entities and one for relationships. For example, the default entity relation would consist of three attributes, one to store the entity identifier, one to store a property name, and one to store a property value. Thus, the table would consist of one tuple for every entity property, where each such tuple has the same value for the id-attribute.

Note that, in the mapping rules **er-rel-8** and **er-rel-13**, we assume for simplicity that both configurations have the same atomic types. Otherwise, we would need to add additional formulas to the head of each rule to define the domain type in the source.

## **Transformation 2** (*E-R-to-relational data transformation*)

mapping **er-rel-7**: (*Every entity relation has a corresponding table.*)

$$\begin{aligned} & \text{c-inst}(\text{B}, \text{table}), \text{d-inst}(\text{B}, \text{T}) \Leftarrow \\ & \text{c-inst}(\text{T}, \text{entityType}), \mathcal{K}_{\text{relation}}^{\text{T}}(\text{R}), \mathcal{K}_{\text{table}}^{\text{T}}(\text{B}). \end{aligned}$$

mapping **er-rel-8**: (*Entities become tuples.*)

$$\begin{aligned} & \text{c-inst}(\text{P}, \text{tuple}), \text{P} \in \text{B}, \text{d-inst}(\text{P}, \text{R}) \Leftarrow \\ & \text{c-inst}(\text{E}, \text{entity}), \text{d-inst}(\text{E}, \text{T}), \mathcal{K}_{\text{tuple}}^{\text{E,T}}(\text{P}), \mathcal{K}_{\text{table}}^{\text{T}}(\text{B}), \mathcal{K}_{\text{relation}}^{\text{T}}(\text{R}). \end{aligned}$$

mapping **er-rel-9**: (*Entity attributes become relation attributes.*)

$$\begin{aligned} & \text{P}[\text{I}]=\text{V}, \text{c-inst}(\text{V}, \text{C}) \Leftarrow \\ & \text{c-inst}(\text{E}, \text{entity}), \text{d-inst}(\text{E}, \text{T}), \text{E}.\text{hasProps:P}_e, \text{P} \in \text{P}_e, \text{P}.\text{hasName:N}, \text{P}.\text{hasVal:V}, \\ & \text{c-inst}(\text{V}, \text{C}), \text{T}.\text{hasAtts:AS}, \mathcal{K}_{\text{attList}}^{\text{AS}}(\text{AL}), \text{rel/AL}[\text{I}]=\text{A}, \text{rel/A}.\text{hasName:N}, \end{aligned}$$

$$\mathcal{K}_{tuple}^{E,T}(P).$$

mapping er-rel-10: (*Create a key value for each entity tuple.*)

$$\begin{aligned} P[I]=V, \text{c-inst}(V, \text{uldString}) \Leftarrow \\ \text{c-inst}(E, \text{entity}), \text{d-inst}(E, T), T.\text{hasAtts:AS}, \mathcal{K}_{attList}^{AS}(\text{AL}), \\ \text{rel/AL}[I]=A, \text{rel/A.hasName:'id'}, \mathcal{K}_{id}^{E,T}(V), \mathcal{K}_{tuple}^{E,T}(P). \end{aligned}$$

mapping er-rel-11: (*Every relationship has a corresponding table.*)

$$\text{c-inst}(B, \text{table}), \text{d-inst}(B, R) \Leftarrow \text{c-inst}(T, \text{relType}), \mathcal{K}_{relation}^T(R), \mathcal{K}_{table}^T(B).$$

mapping er-rel-12: (*Relationships become tuples.*)

$$\begin{aligned} \text{c-inst}(P, \text{tuple}), P \in B, \text{d-inst}(P, R') \Leftarrow \\ \text{c-inst}(R, \text{rel}), \text{d-inst}(R, T), \mathcal{K}_{table}^T(B), \mathcal{K}_{tuple}^{R,T}(P), \mathcal{K}_{relation}^R(R') \end{aligned}$$

mapping er-rel-13: (*Relationship attributes become relation attributes.*)

$$\begin{aligned} P[I]=V, \text{c-inst}(V, C) \Leftarrow \\ \text{c-inst}(R, \text{rel}), \text{d-inst}(R, T), R.\text{hasProps:P}_e, P \in P_e, P.\text{hasName:N}, P.\text{hasVal:V}, \\ \text{c-inst}(V, C), T.\text{hasAtts:AS}, \mathcal{K}_{attList}^{AS}(\text{AL}), \text{rel/AL}[I]=A, \text{rel/A.hasName:N}, \\ \mathcal{K}_{tuple}^{R,T}(P). \end{aligned}$$

mapping er-rel-14: (*Relationship roles become relation attributes.*)

$$\begin{aligned} P[I]=V, \text{c-inst}(V, \text{uldString}) \Leftarrow \\ \text{c-inst}(R, \text{rel}), \text{d-inst}(R, T_r), T_r.\text{hasAtts:AS}, T_r.\text{hasRoles:RL}, R.\text{associates:EL}, \\ \text{EL}[I_e]=E, \text{AS}[I_e]=T_e, \mathcal{K}_{attList}^{AS}(\text{AL}), \mathcal{K}_{attribute}^{R,I_e}(A), \text{rel/AL}[I]=A, \mathcal{K}_{id}^{E,T_e}(V), \mathcal{K}_{tuple}^{R,T_r}(P). \end{aligned}$$

## 4.2.2 XML to Relational Transformations

We consider two XML to relational transformations. The first is given below (Transformation 3) and is called the “edge” approach [77]. This transformation does not consider schema information from the XML configuration and instead transforms the XML document directly to tuples in a pre-defined relational schema.

In the edge approach, a single relational table (called the *edge* table) is used to store an XML document. The relation for the edge table contains five attributes. To illustrate, consider the XML example from Figure 1.2 represented as a tree in Figure 4.1 and in an edge table in Figure 4.2. As shown, each element, attribute, and PCData value is represented as a distinct node. Each node in the XML document, except PCData, is




Figure 4.1: The edge-based, tree representation of the XML document of Figure 1.2. Nodes are annotated with their ordinal number and levels are annotated with their *sourceId* value.

mapped to a tuple in the edge table. The first attribute of the edge table, called *SourceId*, stores an identifier that uniquely identifies the parent node. For example, the root node (*moviedb*) of Figure 4.1 has identifier *id1*, the children of *moviedb* have the identifier *id2*, the children of the first child of *moviedb* have the identifier *id3*, the children of the second child of *moviedb* have the identifier *id5*, and so on. The second attribute of the edge table stores the node's label, which is either an element name or an attribute name (called *Tag* [77]). The third attribute stores the order (*Ordinal*) of a node relative to its parent node. For attributes, the ordinal value is always zero and for the root element the ordinal value is one. The fourth attribute (*TargetId*) stores the *SourceId* for the children of the node or 'NULL' if the node does not have any children. Finally, the fifth attribute (*Data*) stores the PCDATA value, the attribute value, or 'NULL' if neither exists. Note that the *Data* attribute is an example of inlining [74, 77].

The mapping rules **edge-1** through **edge-6** below create the pre-defined relational schema for the edge transformation. The last mapping rule (**edge-7**) uses the *findEdge* intensional predicate to generate tuples, which are stored in the edge schema.

Our motivation for including the XML edge-based mapping is to highlight the flexibility of the ULD transformation language. In particular, the transformation expressed in the ULD works regardless of whether a schema is defined for the XML configuration. Existing meta-data-model approaches are based on the assumption that schema is always defined

*EdgeTable*

| <i>SourceId</i> | <i>Tag</i> | <i>Ordinal</i> | <i>TargetId</i> | <i>Data</i>          |
|-----------------|------------|----------------|-----------------|----------------------|
| id1             | moviedb    | 1              | id2             | 'NULL'               |
| id2             | movie      | 1              | id3             | 'NULL'               |
| id2             | movie      | 2              | id5             | 'NULL'               |
| id3             | title      | 1              | 'NULL'          | 'The Usual Suspects' |
| id3             | studio     | 2              | 'NULL'          | 'Gramercy'           |
| id3             | genre      | 3              | 'NULL'          | 'Thriller'           |
| id3             | actor      | 4              | id4             | 'Spacey, Kevin'      |
| id4             | role       | 0              | 'NULL'          | 'Supporting Actor'   |
| id5             | title      | 1              | 'NULL'          | 'Meet the Parents'   |
| id5             | studio     | 2              | 'NULL'          | 'Universal'          |
| id5             | genre      | 3              | 'NULL'          | 'Comedy'             |
| id5             | actor      | 4              | id6             | 'De Niro, Robert'    |
| id6             | role       | 0              | 'NULL'          | 'Leading Actor'      |

Figure 4.2: The result of applying the edge transformation.

for a source. Thus, they cannot express such data-only transformations because access to data can occur only through schema structures.

Finally, the edge-based transformation maps XML data directly to relational data using an agreed-upon convention. We note that, in general, conventions are often introduced in transformations among heterogeneous information sources due to the differences in the constructs of the data models. In the edge approach, XML information is stored in the relational data model using a predefined table to represent an ordered tree, because the relational data model does not have constructs for explicitly representing hierarchical information.

**Transformation 3** (*An XML-to-relational edge transformation*)

mapping edge-1: (*Construct the edge table.*)

```
c-inst(edgeTbl, relation), edgeTbl.hasName:'EdgeTable', c-inst('EdgeTable', uldString),
c-inst(edgeTblAtts, attList), edgeTbl.hasAtts:edgeTblAtts  $\Leftarrow$  .
```

mapping edge-2: (*Construct the source id attribute.*)

```
c-inst(sourceId, attribute), sourceId.hasName:'SourceId', c-inst('SourceId', uldString),
sourceId.hasDomain:'uldString', edgeTblAtts[1]=sourceId  $\Leftarrow$  .
```

mapping edge-3: (*Construct the tag attribute.*)

```
c-inst(tag, attribute), tag.hasName:'Tag', c-inst('Tag', uldString),
```

tag.hasDomain:'uldString', edgeTblAtts[2]=tag  $\Leftarrow$  .

mapping edge-4: (*Construct the ordinal attribute.*)

c-inst(ordinal, attribute), ordinal.hasName:'Ordinal', c-inst('Ordinal', uldString),  
ordinal.hasDomain:'uldInteger', edgeTblAtts[3]=ordinal  $\Leftarrow$  .

mapping edge-5: (*Construct the target id attribute.*)

c-inst(targetId, attribute), targetid.hasName:'TargetId', c-inst('TargetId', uldString),  
targetId.hasDomain:'uldString', edgeTblAtts[4]=targetId  $\Leftarrow$  .

mapping edge-6: (*Construct the data attribute.*)

c-inst(data, attribute), data.hasName:'Data', c-inst('Data', uldString),  
data.hasDomain:'uldString', edgeTblAtts[5]=data  $\Leftarrow$  .

mapping edge-7: (*Convert elements to tuples.*)

c-inst(ID, tuple), tuple[1]=S, tuple[2]=T, tuple[3]=O, tuple[4]=R, tuple[5]=D,  
c-inst(S, uldString), c-inst(T, uldString), c-inst(O, uldInteger), c-inst(R, uldString),  
c-inst(D, uldString), tuple $\in$ edgeTbl  $\Leftarrow$   
findEdge(ID, S, T, O, R, D).

% root node without child element and text content

findEdge(ID, ID, T, 1, 'NULL', 'NULL')  $\Leftarrow$  c-inst(D, doc), D.hasRoot:E, E.hasTag:T,  
 $\neg$ hasSubElemsOrAtts(E), E.hasChildren:C,  $\neg$ hasPCData(C),  $\mathcal{K}_{tuple}^E$ (ID).

% root node with text content, but without child element

findEdge(ID, ID, T, 1, 'NULL', D)  $\Leftarrow$  c-inst(D, doc), D.hasRoot:E, E.hasTag:T,  
 $\neg$ hasSubElemsOrAtts(E), E.hasChildren:C, hasPCData(C, D),  $\mathcal{K}_{tuple}^E$ (ID).

% root node with child element, but without text content

findEdge(ID, ID, T, 1, R, 'NULL')  $\Leftarrow$  c-inst(D, doc), D.hasRoot:E, E.hasTag:T,  
hasSubElemsOrAtts(E),  $\mathcal{K}_{children}^E$ (R), E.hasChildren:C,  $\neg$ hasPCData(C),  
 $\mathcal{K}_{tuple}^E$ (ID).

% root node with child element and text content

findEdge(ID, ID, T, 1, R,D)  $\Leftarrow$  c-inst(D, doc), D.hasRoot:E, E.hasTag:T,  
hasSubElemsOrAtts(E),  $\mathcal{K}_{children}^E$ (R), E.hasChildren:C, hasPCData(C, D),  
 $\mathcal{K}_{tuple}^E$ (ID).

% attribute

findEdge(ID, S, T, 0, 'NULL', D)  $\Leftarrow$  c-inst(E, element),  $\mathcal{K}_{children}^E$ (S), E.hasAtts:AS,  
A $\in$ AS, A.hasName:T, A.hasVal:D,  $\mathcal{K}_{tuple}^A$ (ID)

% element without child element and text content

findEdge(ID, S, T, O, 'NULL', 'NULL')  $\Leftarrow$  c-inst( $E_p$ , element),  $E_p$ .hasChildren:C,

```

    hasChild(C, E, O),  $\mathcal{K}_{children}^{E_p}(S)$ , E.hasName:T,  $\neg$ hasSubElemsOrAtts(E),
    E.hasChildren:C,  $\neg$ hasPCData(C),  $\mathcal{K}_{tuple}^E(ID)$ .
% element with text content, but without child element
findEdge(ID, S, T, O, 'NULL', D)  $\leftarrow$  c-inst( $E_p$ , element),  $E_p$ .hasChildren:C,
    hasChild(C, E, O),  $\mathcal{K}_{children}^{E_p}(S)$ , E.hasName:T,  $\neg$ hasSubElemsOrAtts(E),
    E.hasChildren:C, hasPCData(C, D),  $\mathcal{K}_{tuple}^E(ID)$ .
% element with child element, but without text content
findEdge(ID, S, T, O, R, 'NULL')  $\leftarrow$  c-inst( $E_p$ , element),  $E_p$ .hasChildren:C,
    hasChild(C, E, O),  $\mathcal{K}_{children}^{E_p}(S)$ , E.hasName:T, hasSubElemsOrAtts(E),  $\mathcal{K}_{children}^E(R)$ ,
    E.hasChildren:C,  $\neg$ hasPCData(C),  $\mathcal{K}_{tuple}^E(ID)$ .
% element with child element and text content
findEdge(ID, S, T, O, R, D)  $\leftarrow$  c-inst( $E_p$ , element),  $E_p$ .hasChildren:C, hasChild(C, E, O),
     $\mathcal{K}_{children}^{E_p}(S)$ , E.hasName:T, hasSubElemsOrAtts(E),  $\mathcal{K}_{children}^E(R)$ , E.hasChildren:C,
    hasPCData(C, D),  $\mathcal{K}_{tuple}^E(ID)$ .

hasPCData(C, D)  $\leftarrow$  c-inst(C, fullContent), C.hasNode:D, c-inst(D, pcdata).
hasPCData(C, D)  $\leftarrow$  c-inst(C, fullContent), C.hasContent:C', hasPCData(C', D).

hasPCData(C)  $\leftarrow$  hasPCData(C, D).

hasSubElems(C)  $\leftarrow$  c-inst(C, fullContent), C.hasNode:N, c-inst(N, element).
hasSubElems(C)  $\leftarrow$  c-inst(C, fullContent), C.hasChildren:C', hasSubElems(C').

hasSubElemsOrAtts(E)  $\leftarrow$  c-inst(E, element), E.hasAtts:AS, |AS| $\neq$ 0.
hasSubElemsOrAtts(E)  $\leftarrow$  c-inst(E, element), E.hasChildren:C, hasSubElems(C).

hasChild(C, E, 1)  $\leftarrow$  c-inst(C, fullContent), C.hasNode:E, c-inst(E, element).
hasChild(C, E, O)  $\leftarrow$  c-inst(C, fullContent), C.hasContent:C', hasChild(C', E, O'),
    O=O'+1.

```

In contrast to the edge transformation, the “attribute” approach [74, 77] uses schema information to map an XML document to a relational database. The attribute transformation (Transformation 4) is given below. Figure 4.3 shows the result of applying the attribute transformation on the XML DTD and instance of Figure 1.2. The transformation maps each element type to a relation consisting of an id-attribute. Each element-type that is in the content specification of another element type forms a new relation, which stores the sub-element connection between elements. This approach is similar to the mapping of relationship types to relations in the E-R to relational transformation. The attribute



c-inst(T, elemType), T.hasAtts:AD,  $\mathcal{K}_{attList}^{AD}(AL)$ ,  $\mathcal{K}_{idatt}^T(A)$ , .

mapping att-3: (*Element type attributes become relation attributes.*)

c-inst( $A_r$ , attribute),  $A_r \in AL$ ,  $A_r$ .hasName:N, c-inst(N, uldString),  
 $A_r$ .hasDomain:'uldString'  $\Leftarrow$   
 c-inst(T, elemType), T.hasAtts:AD,  $A_e \in AD$ ,  $A_e$ .hasName:N,  $\mathcal{K}_{attribute}^{A_e}(A_r)$ ,  
 $\mathcal{K}_{attList}^{AD}(AL)$ .

mapping att-4: (*Element types with content have additional relation attributes.*)

c-inst(A, attribute),  $A \in AL$ , A.hasName:'content', c-inst('content', uldString),  
 A.hasDomain:'uldString'  $\Leftarrow$   
 c-inst(T, elemType), T.hasModel:CS, c-inst(CS, simpleSpec), CS.isPCData:true,  
 $\mathcal{K}_{content}^T(A)$ , T.hasAtts:AD,  $\mathcal{K}_{attList}^{AD}(AL)$ .

mapping att-5: (*Element-type subelements become edge relations.*)

c-inst(R, relation), R.hasName:N, c-inst(N, uldString), R.hasAtts:AL,  
 c-inst(AL, attList)  $\Leftarrow$   
 c-inst( $T_1$ , elemType),  $T_1$ .hasModel:CS, isSubElement(CS,  $T_2$ ),  $T_1$ .hasName: $N_1$ ,  
 $T_2$ .hasName: $N_2$ ,  $N = N_1 \cdot N_2$ ,  $\mathcal{K}_{relation}^{T_1, T_2}(R)$ ,  $\mathcal{K}_{attList}^{T_1, T_2}(AL)$ .

mapping att-6: (*Subelement edge relations have two identifying attributes.*)

c-inst( $A_1$ , attribute), c-inst( $A_2$ , attribute),  $AL[1]=A_1$ ,  $AL[2]=A_2$ ,  $A_1$ .hasName: $N_1$ ,  
 c-inst( $N_1$ , uldString),  $A_2$ .hasName: $N_2$ , c-inst( $N_2$ , uldString)  $\Leftarrow$   
 c-inst( $T_1$ , elemType),  $T_1$ .hasModel:CS, isSubElemType(CS,  $T_2$ ),  $T_1$ .hasName: $N_1$ ,  
 $T_2$ .hasName: $N_2$ ,  $\mathcal{K}_{relation}^{T_1, T_2}(R)$ ,  $\mathcal{K}_{attList}^{T_1, T_2}(AL)$ ,  $\mathcal{K}_{att1}^{T_1, T_2}(A_1)$ ,  $\mathcal{K}_{att2}^{T_1, T_2}(A_2)$ .

mapping att-7: (*Create a table for each element type.*)

c-inst(B, table), d-inst(B, R)  $\Leftarrow$  c-inst(T, elemType),  $\mathcal{K}_{relation}^T(R)$ ,  $\mathcal{K}_{table}^T(B)$ .

mapping att-8: (*Create a table for each subelement type.*)

c-inst(B, table), d-inst(B, R)  $\Leftarrow$   
 c-inst( $T_1$ , elemType),  $T_1$ .hasModel:CS, isSubElemType(CS,  $T_2$ ),  $\mathcal{K}_{relation}^{T_1, T_2}(R)$ ,  
 $\mathcal{K}_{table}^{T_1, T_2}(B)$ .

mapping att-9: (*Create a tuple for each element.*)

c-inst(P, tuple),  $P \in B \Leftarrow$  c-inst(E, element), d-inst(E, T),  $\mathcal{K}_{table}^T(B)$ ,  $\mathcal{K}_{tuple}^E(P)$ .

mapping att-10: (*Create an identifying-attribute value for each element.*)

$P[I]=V$ , c-inst(V, uldString)  $\Leftarrow$

$c\text{-inst}(E, \text{element}), d\text{-inst}(E, T), \mathcal{K}_{table}^T(R), \mathcal{K}_{idatt}^T(A),$   
 $rel/R.\text{hasAtts}:AL, rel/AL[I]=A, \mathcal{K}_{tuple}^E(P), \mathcal{K}_{idatt}^E(V).$

mapping att-11: (*Create an attribute value for each element attribute.*)

$P[I]=V, c\text{-inst}(V, \text{uldString}) \Leftarrow$   
 $c\text{-inst}(E, \text{element}), d\text{-inst}(E, T), T.\text{hasAtts}:AD, A_e \in AD, A_e.\text{hasVal}:V,$   
 $\mathcal{K}_{attribute}^{A_e}(A_r), \mathcal{K}_{attList}^{AD}(AL), rel/AL[I]=A_r, \mathcal{K}_{tuple}^E(P).$

mapping att-12: (*Create an attribute value for elements with only text content.*)

$P[I]=V, c\text{-inst}(V, \text{uldString}) \Leftarrow$   
 $c\text{-inst}(E, \text{element}), d\text{-inst}(E, T), T.\text{hasAtts}:AD, T.\text{hasModel}:CS,$   
 $c\text{-inst}(CS, \text{simpleSpec}), CS.\text{isPCData}:true, E.\text{hasChildren}:C, C.\text{hasNode}:V,$   
 $c\text{-inst}(V, \text{pcdata}), \mathcal{K}_{content}^T(A), \mathcal{K}_{attList}^{AD}(AL), rel/AL[I]=A, \mathcal{K}_{tuple}^E(P).$

mapping att-13: (*Create a tuple for each subelement.*)

$c\text{-inst}(P, \text{tuple}), P[1]=V_1, P[2]=V_2, P \in B \Leftarrow$   
 $c\text{-inst}(E_1, \text{element}), d\text{-inst}(E_1, T_1), E_1.\text{hasChildren}:C, \text{isSubElement}(C, E_2),$   
 $d\text{-inst}(E_2, T_2), \mathcal{K}_{table}^{T_1, T_2}(B), \mathcal{K}_{tuple}^{E_1, E_2}(P), \mathcal{K}_{idatt}^{E_1}(V_1), \mathcal{K}_{idatt}^{E_2}(V_2).$

$\text{isSubElemType}(CS, T) \leftarrow c\text{-inst}(CS, \text{mixedSpec}), T \in CS.$   
 $\text{isSubElemType}(CS, T) \leftarrow c\text{-inst}(CS, \text{elemRE}), CS.\text{hasType}:T.$   
 $\text{isSubElemType}(CS, T) \leftarrow c\text{-inst}(CS, \text{regExp}), CS.\text{hasRE}:CS', \text{isSubElemType}(CS', T).$   
 $\text{isSubElemType}(CS, T) \leftarrow c\text{-inst}(CS, \text{regExp}), CS.\text{hasTail}:CS', \text{isSubElemType}(CS', T).$   
 $\text{isSubElement}(C, E) \leftarrow c\text{-inst}(C, \text{fullContent}), C.\text{hasNode}:E, c\text{-inst}(E, \text{element}).$   
 $\text{isSubElement}(C, E) \leftarrow c\text{-inst}(C, \text{fullContent}), C.\text{hasContent}:C', \text{isSubElement}(C', E).$

### 4.2.3 Schema Transformation

The previous transformations perform data-model-to-data-model conversions. Here, we consider using the ULD transformation language to convert data among particular schemas, where each schema is represented in a different data model.

The following transformation converts the XML source given in Figure 1.2 to a valid RDF source that conforms to the RDF Schema of Figure 1.3. The transformation must consider schematic differences, for example, genre information is represented as data in the XML example and as schema in the RDF example, as well as format differences, for example, names are represented in slightly different ways between the schemas. We make

the following assumptions concerning the transformation. The function *toRDFName* takes a name formatted last-name first and returns a new name formatted first-name first. The identifiers *film*, *thriller*, *comedy*, and so on are defined in the RDF Schema. Similarly, the identifiers *movie*, *genre*, and so on are defined in the XML DTD. And finally, the Skolem function  $\mathcal{K}_{id}^M(U)^{uldURI}$  unifies the variable U with a unique and valid URI value. (Note that *uldURI* is another modifier to the  $\mathcal{K}$  function that in general, requests identifiers that conform to a specific atomic type such as strings, integers, URIs, and so on.)

The transformation below is a schema-to-schema mapping for schemas represented in distinct data models. As a result of the mapping, XML data is directly converted to valid RDF data. Alternatively, in schema-mapping approaches, the transformation below must typically be expressed in a common data model. The XML DTD and RDF Schema would first be converted to the common data model (which in general is not trivial). The schema mapping is then defined in the common data model. And finally, the transformation is implemented by converting the mapping rules to an executable program, or by converting source data to the common data model, executing the transformation within the common data model, and then converting the resulting data to the target data model. As shown below, the flexibility of the ULD allows schema mappings to be expressed directly, without the need for additional mappings and conversions to a common data model.

**Transformation 5** (*A DTD-to-RDF-Schema transformation*)

mapping schema-1: (*Movie elements become film resources.*)

$$\text{c-inst}(\mathbf{R}, \text{simpleRes}), \text{d-inst}(\mathbf{R}, \text{film}), \mathbf{R}.\text{hasURI:U}, \text{c-inst}(\mathbf{U}, \text{uldURI}) \Leftarrow \\ \text{d-inst}(\mathbf{M}, \text{movie}), \mathcal{K}_{resource}^M(\mathbf{R}), \mathcal{K}_{id}^M(\mathbf{U})^{uldURI}.$$

mapping schema-2: (*Thriller movies become thriller-film resources.*)

$$\text{d-inst}(\mathbf{R}, \text{thriller}) \Leftarrow \\ \text{d-inst}(\mathbf{M}, \text{movie}), \mathbf{M}.\text{hasChildren:C}, \text{isSubElement}(\mathbf{C}, \mathbf{M}'), \text{d-inst}(\mathbf{M}', \text{genre}), \\ \mathbf{M}'.\text{hasChildren:C'}, \text{hasPCData}(\mathbf{C}', \text{'Thriller'}), \mathcal{K}_{resource}^M(\mathbf{R}).$$

mapping schema-3: (*Comedy movies become thriller-film resources.*)

$$\text{d-inst}(\mathbf{R}, \text{comedy}) \Leftarrow \\ \text{d-inst}(\mathbf{M}, \text{movie}), \mathbf{M}.\text{hasChildren:C}, \text{isSubElement}(\mathbf{C}, \mathbf{M}'), \text{d-inst}(\mathbf{M}', \text{genre}), \\ \mathbf{M}'.\text{hasChildren:C'}, \text{hasPCData}(\mathbf{C}', \text{'Comedy'}), \mathcal{K}_{resource}^M(\mathbf{R}).$$

mapping schema-4: (*Title elements become film title triples.*)

$$\begin{aligned} & \text{c-inst}(V, \text{literal}), \text{c-inst}(T, \text{triple}), T.\text{hasPred}:\text{title}, T.\text{hasSubj}:R, T.\text{hasObj}:V \Leftarrow \\ & \text{d-inst}(M, \text{movie}), M.\text{hasChildren}:C, \text{isSubElement}(C, M'), \text{d-inst}(M', \text{title}), \\ & M'.\text{hasChildren}:C', \text{hasPCData}(C', V), \mathcal{K}_{resource}^M(R), \mathcal{K}_{title}^M(T). \end{aligned}$$

mapping schema-5: (*Actor elements become actor resources.*)

$$\begin{aligned} & \text{c-inst}(R, \text{simpleRes}), \text{d-inst}(R, \text{actor}), R.\text{hasURI}:U, \text{c-inst}(U, \text{uldURI}) \Leftarrow \\ & \text{d-inst}(A, \text{actor}), \mathcal{K}_{resource}^A(R), \mathcal{K}_{id}^A(U)^{\text{uldURI}}. \end{aligned}$$

mapping schema-6: (*Actor names become hasName properties.*)

$$\begin{aligned} & \text{c-inst}(V, \text{literal}), \text{c-inst}(T, \text{triple}), T.\text{hasPred}:\text{hasName}, T.\text{hasSubj}:R, T.\text{hasObj}:V \Leftarrow \\ & \text{d-inst}(A, \text{actor}), A.\text{hasChildren}:C, \text{hasPCData}(C, N), V=\text{toRDFName}(N), \\ & \mathcal{K}_{resource}^A(R), \mathcal{K}_{name}^A(T). \end{aligned}$$

mapping schema-7: (*Actor subelements are connected to film resources.*)

$$\begin{aligned} & \text{c-inst}(T, \text{triple}), T.\text{hasPred}:\text{involved}, T.\text{hasSubj}:R_m, T.\text{hasObj}:R_a \Leftarrow \\ & \text{d-inst}(M, \text{movie}), M.\text{hasChildren}:C, \text{isSubElement}(C, A), \text{d-inst}(A, \text{actor}), \\ & \mathcal{K}_{resource}^M(R_m), \mathcal{K}_{resource}^A(R_a), \mathcal{K}_{involved}^{M,A}(T). \end{aligned}$$

$$\text{isSubElement}(C, E) \leftarrow \text{c-inst}(C, \text{fullContent}), C.\text{hasNode}:E, \text{c-inst}(E, \text{element}).$$

$$\text{isSubElement}(C, E) \leftarrow \text{c-inst}(C, \text{fullContent}), C.\text{hasContent}:C', \text{isSubElement}(C', E).$$

$$\text{hasPCData}(C, D) \leftarrow \text{c-inst}(C, \text{fullContent}), C.\text{hasNode}:D, \text{c-inst}(D, \text{pcdata}).$$

$$\text{hasPCData}(C, D) \leftarrow \text{c-inst}(C, \text{fullContent}), C.\text{hasContent}:C', \text{hasPCData}(C', D).$$

#### 4.2.4 Transformations for Standard Serializations

The last transformation (Transformation 6) of this section converts RDF data directly into XML data. Figure 4.4 gives the result of applying the mapping to the RDF configuration shown in Figure 1.3. The resulting XML document is similar to the standard RDF serialization in XML [47]. A serialization is, fundamentally, a data-model mapping in which one data model is used as the interchange format for information expressed in another data model. One of the most widely used data models for serialization is XML. For example, both RDF and Topic Maps have multiple XML serializations.

```

<rdf>
  <description about="#m1">
    <title>The Usual Suspects</title>
    <involved resource="#ks"/>
    <won resource="#a1"/>
  </description>
  <description about="#ks">
    <hasName>Kevin Space</hasName>
  </description>
  <description about="#a1">
    <recipient resource="#ks"/>
    <category>Supporting Actor</category>
  </description>
</rdf>

```

Figure 4.4: The result of applying the RDF serialization mapping.

We note that it is not possible to define an XML DTD for the standard RDF serialization. In particular, the element tags present in an XML document depend on the property names used in the RDF source. Therefore, one benefit of using the ULD transformation language is that it provides a formal, explicit, and executable definition of the serialization. In addition, serializations often use a number of conventions and mix data and schema levels. For example, RDF schema information and instance data are both mapped to XML elements and attributes (that is, instance data) in the standard RDF serialization. The flexibility of the ULD enables the specification of these various styles of mappings that mix data-model, schema, and data levels.

**Transformation 6** (*An RDF-to-XML serialization*)

mapping rdxml-1: (*Generate a document and an RDF root tag.*)

c-inst(rdfDoc, doc), c-inst(rdfRoot, element), rdfRoot.hasTag:'rdf',  
 rdfRoot.hasAtts:rdfRootAtts, c-inst('rdf', uldString),  
 c-inst(rdfRootAtts, emptyContent)  $\Leftarrow$  .

mapping rdxml-2: (*Each resource that is the subject of a triple generates an element.*)

c-inst(E, element), E.hasTag:'description', c-inst('description', uldString), E.hasAtts:AS,  
 c-inst(AS, attSet)  $\Leftarrow$   
 c-inst(T, triple), T.hasSubj:S,  $\mathcal{K}_{desc}^S(E)$ ,  $\mathcal{K}_{attList}^S(AS)$ .

mapping rdxml-3: (*Add first subelement of the root.*)

rdxRoot.hasChildren:rdxRootContent, c-inst(rdxRootContent, fullContent),  
rdxRootContent.hasNode:E  $\Leftarrow$   
c-inst(T, triple), T.hasSubj:S,  $\mathcal{K}_{desc}^S(E)$ ,  $xml/\neg$ c-inst(rdxRootContent, fullContent).

mapping rdxml-4: (*Add remaining subelements of the root.*)

C.hasContent:C', c-inst(C', fullContent), C'.hasNode:E  $\Leftarrow$   
c-inst(T, triple), T.hasSubj:S,  $\mathcal{K}_{desc}^S(E)$ , findLastContent(rdxRootContent, C),  
 $\mathcal{K}_{content}^S(C')$ .

mapping rdxml-5: (*If every subject has been converted, add empty content.*)

C.hasContent:C', c-inst(C', emptyContent)  $\Leftarrow$   
 $\neg$ existsSubjNotConverted(rdxRootContent), findLastContent(rdxRootContent, C),  
 $\mathcal{K}_{emptycont}^S(C')$ .

mapping rdxml-6: (*Add the subject resource URI as an attribute.*)

E.hasAtts:AS, c-inst(AS, attSet), A $\in$ AS, c-inst(A, attribute), A.hasName:'about',  
c-inst('about', uldString), A.hasVal:U, c-inst(U, cdata)  $\Leftarrow$   
c-inst(T, triple), T.hasSubj:S, S.hasURI:U,  $\mathcal{K}_{desc}^S(E)$ ,  $\mathcal{K}_{attSet}^S(AS)$ ,  $\mathcal{K}_{attribute}^S(A)$ .

mapping rdxml-7: (*Generate first subelement of subject element.*)

E.hasChildren:C, c-inst(C, fullContent), C.hasNode:E', c-inst(E', element), E'.hasTag:U,  
E'.hasAtts:AS, c-inst(AS, attSet), E'.hasContent:C<sub>p</sub>, c-inst(C<sub>p</sub>, fullContent), C<sub>p</sub>.hasNode:O,  
c-inst(O, pcdData), C<sub>p</sub>.hasContent:C<sub>e</sub>, c-inst(C<sub>e</sub>, emptyContent)  $\Leftarrow$   
c-inst(T, triple), T.hasSubj:S, T.hasPred:P, P.hasURI:U, T.hasObj:O,  
c-inst(O, literal),  $\mathcal{K}_{desc}^S(E)$ ,  $\mathcal{K}_{children}^S(C)$ ,  $xml/\neg$ c-inst(C, fullContent),  $\mathcal{K}_{element}^{S,P,O}(E')$ ,  
 $\mathcal{K}_{predicate}^{S,P,O}(C_p)$ ,  $\mathcal{K}_{empty}^{S,P,O}(C_e)$ ,  $\mathcal{K}_{attSet}^{S,P,O}(AS)$ .

mapping rdxml-8: (*Generate first subelement of subject element.*)

E.hasChildren:C, c-inst(C, fullContent), C.hasNode:E', c-inst(E', element), E'.hasTag:U<sub>p</sub>,  
E'.hasAtts:AS, c-inst(AS, attSet), A $\in$ AS, c-inst(A, attribute), A.hasName:'resource',  
c-inst('resource', uldString), A.hasVal:U<sub>o</sub>, c-inst(U<sub>o</sub>, cdata), E'.hasContent:C',  
c-inst(C', emptyContent)  $\Leftarrow$   
c-inst(T, triple), T.hasSubj:S, T.hasPred:P, P.hasURI:U<sub>p</sub>, T.hasObj:O, O.hasURI:U<sub>o</sub>,  
 $\mathcal{K}_{desc}^S(E)$ ,  $\mathcal{K}_{children}^S(C)$ ,  $xml/\neg$ c-inst(C, fullContent),  $\mathcal{K}_{element}^{S,P,O}(E')$ ,  $\mathcal{K}_{empty}^{S,P,O}(C')$ ,  
 $\mathcal{K}_{attSet}^{S,P,O}(AS)$ ,  $\mathcal{K}_{object}^{S,P,O}(A)$ .

mapping rdxml-9: (*Generate remaining subelements of subject element.*)

C.hasContent:C', c-inst(C', fullContent), C'.hasNode:E', c-inst(E', element), E'.hasTag:U,

$E'.hasAtts:AS, c-inst(AS, attSet), E'.hasContent:C_p, c-inst(C_p, fullContent), C_p.hasNode:O,$   
 $c-inst(O, pcdData), C_p.hasContent:C_e, c-inst(C_e, emptyContent) \Leftarrow$   
 $c-inst(T, triple), T.hasSubj:S, T.hasPred:P, P.hasURI:U, T.hasObj:O,$   
 $c-inst(O, literal), \mathcal{K}_{desc}^S(E), \mathcal{K}_{children}^S(C), findLastContent(C, C'), \mathcal{K}_{element}^{S,P,O}(E'),$   
 $\mathcal{K}_{predicate}^{S,P,O}(C_p), \mathcal{K}_{empty}^{S,P,O}(C_e), \mathcal{K}_{attSet}^{S,P,O}(AS).$

mapping rdxml-10: (*Generate remaining subelements of subject element.*)

$C.hasContent:C', c-inst(C', fullContent), C'.hasNode:E', c-inst(E', element), E'.hasTag:U_p,$   
 $E'.hasAtts:AS, c-inst(AS, attSet), A \in AS, c-inst(A, attribute), A.hasName:'resource',$   
 $c-inst('resource', uldString), A.hasVal:U_o, c-inst(U_o, cdata), E'.hasContent:C',$   
 $c-inst(C', emptyContent) \Leftarrow$   
 $c-inst(T, triple), T.hasSubj:S, T.hasPred:P, P.hasURI:U_p, T.hasObj:O, O.hasURI:U_o,$   
 $\mathcal{K}_{desc}^S(E), \mathcal{K}_{children}^S(C), findLastContent(C, C'), \mathcal{K}_{element}^{S,P,O}(E'), \mathcal{K}_{empty}^{S,P,O}(C'),$   
 $\mathcal{K}_{attSet}^{S,P,O}(AS), \mathcal{K}_{object}^{S,P,O}(A).$

mapping rdxml-11: (*If every triple has been converted, add empty content.*)

$C'.hasContent:C_e, c-inst(C_e, emptyContent) \Leftarrow$   
 $c-inst(T, triple), T.hasSubj:S, \neg existsTripleNotConverted(S), \mathcal{K}_{desc}^S(E),$   
 $xml/E.hasContent:C, findLastContent(C, C'), \mathcal{K}_{emptycont}^S(C_e).$

$existsSubjNotConverted(C) \leftarrow xml/c-inst(C, fullContent), c-inst(T, triple), T.hasSubj:S,$   
 $\mathcal{K}_{desc}^S(E), \neg isSubElement(C, E).$

$existsTripleNotConverted(S) \leftarrow c-inst(T, triple), T.hasSubj:S, \mathcal{K}_{desc}^S(E), E.hasContent:C,$   
 $T.hasPred:P, T.hasObj:O, \mathcal{K}_{element}^{S,P,O}(E'), \neg isSubElement(C, E').$

$isSubElement(C, E) \leftarrow xml/c-inst(C, fullContent), xml/C.hasNode:E,$   
 $xml/c-inst(E, element).$

$isSubElement(C, E) \leftarrow xml/c-inst(C, fullContent), xml/C.hasContent:C',$   
 $xml/isSubElement(C', E).$

$findLastContent(C_1, C_1) \leftarrow xml/c-inst(C_1, fullContent), \neg hasFilledContent(C_1).$

$findLastContent(C_1, C_3) \leftarrow xml/c-inst(C_1, fullContent), xml/C_1.hasContent:C_2,$   
 $xml/\neg c-inst(C_2, emptyContent), findLastContent(C_2, C_3).$

$hasFilledContent(C) \leftarrow xml/c-inst(C, fullContent), xml/C.hasContent:C'.$

## 4.3 Analyzing ULD Transformations

We consider two applications of reasoning for ULD mapping rules. By reasoning, we mean analyzing mapping rules prior to their execution using ULD axioms and data-model constraints. The first application we consider ensures that a mapping rule is *correct*, that is, the resulting target configuration does not violate the ULD axioms and data-model constraints. The second application of reasoning is directed at helping a user define a transformation. In particular, we want to exploit the ULD axioms and data-model constraints to “fill-out” incomplete mapping rules. As an ambitious example, a user may specify that E-R entity types should be mapped to relations in the relational model without giving the details of such a mapping, where the rest of the mapping is filled in by a reasoning system. (Note that similar approaches exist for schema mappings [68] and we consider only simple cases of the more general problem in this chapter.)

This section focuses on ensuring mapping rules are correct with respect to ULD axioms and briefly discusses the use of data-model constraints for both ensuring correctness and elaborating partial mapping rules.

### 4.3.1 Well-Typed Transformations

Here we describe a structural approach for verifying that a transformation is *well-typed*. Based on the axioms of the ULD (from Chapter 2), we define a set of rules to annotate variables used in mapping rules with type information, and then use this type information to verify that a transformation is well-formed (with respect to the types).

Our overriding goal is to ensure that a transformation will result in a consistent target configuration and notify the user if the transformation can produce an inconsistent result (or in general, if a mapping rule is not well-formed). We only consider individual mapping rules. In general, a transformation should be considered as a whole; each individual mapping rule can be consistent, but when taken together, they may not produce a consistent transformation.<sup>4</sup>

---

<sup>4</sup>Note that the converse is not true, that is, if a mapping rule can produce an inconsistent configuration, it is not possible to construct another mapping rule to make the configuration consistent. Intuitively, a configuration is inconsistent if it asserts a fact that violates a ULD or data-model constraint. Therefore,

A mapping rule is inconsistent if it can produce target facts that violate ULD axioms or data-model constraints. We also consider whether the body of a mapping rule can be satisfied, that is, whether it is possible to construct a valid configuration that produces a variable assignment for the rule (or for rule bodies that do not contain variables, whether they are false for every possible configuration). No updates will result from executing a mapping rule with a rule body that is not satisfiable. This mapping rule is not considered well-formed even though it does not produce an inconsistent target configuration.

Checking for well-typed mapping rules does not guarantee that a mapping rule produces a consistent target configuration, only that the result is consistent with respect to typing constraints. For example, consider the following mapping rule (Example 11), which maps E-R attributes to relations in the relational data model.

**Example 11** (*A well-typed transformation rule*)

$$\begin{array}{l} \text{c-inst}(\text{T}, \text{relation}), \text{T.hasName:N}, \text{c-inst}(\text{N}, \text{uldString}) \Leftarrow \\ \text{c-inst}(\text{A}, \text{attribute}), \text{A.hasName:N}, \mathcal{K}^A(\text{T}). \end{array}$$

This rule is considered well-typed (which we define below) and thus is valid with respect to typing. However, the target configuration generated from the rule may not be consistent. In particular, multiple attributes with the same name are allowed in the source configuration (which uses an E-R data model), however, each relation in the target configuration (a relational data model) must contain a unique name. Therefore, if there is more than one attribute with the same name in the source, the resulting target configuration will be inconsistent with respect to the relational data-model constraint (hasName keyOf relation).

The *type annotation* rules in Figures 4.5, 4.6, and 4.7 are used to determine the types of variables and constants in a mapping rule. If the antecedent of a rule holds (the expression above the line), the type annotation in the consequent of the rule is asserted (the expression below the line). A type annotation takes the form  $c/v :: t$ , which states that  $v$ , a variable

---

such a fact would need to be removed from the configuration to make it consistent, and it is not possible to remove facts from a configuration using the ULD transformation language.

or constant in the mapping rule, has type  $t$  in configuration  $c$ . An expression such as  $c/c\text{-inst}(x, p)$  (for  $x$  and  $p$  constants or variables and  $c$  a configuration name) in an annotation rule is matched against a similar expression in a mapping rule. (Note that we assume the default source and target configurations if the configuration name is not given, as in the previous example.) For example, the formula  $rel/c\text{-inst}(T, \text{relation})$  in the previous mapping rule (Example 11) would match the expression  $c/c\text{-inst}(x, p)$  in a type annotation rule, where  $x$  is unified with  $T$ ,  $p$  is unified with the constant ‘relation’, and  $c$  is unified with the (default) configuration  $rel$ .

Type annotation rules may also use data-model information from a configuration to annotate variables and constants. Data-model information is accessed using expressions of the form  $c \models b$ , which are true if the variables of a formula  $b$  can be unified in configuration  $c$ . The expression  $c \models b$  is essentially a ULD query, which finds a variable assignment for  $b$  in configuration  $c$  through unification. Data-model expressions can require certain terms to be ground prior to their application. For example, the expression  $c \models \text{ct-inst}(p^*, s)$  requires  $p$  to be either a constant, or a variable unified with a constant.

Type annotation rules are repeatedly applied to each formula in a mapping rule until no additional types are generated. That is, annotation rules are applied until a fixpoint is reached. We note that this approach to typing is similar to the system of Mycroft and O’Keefe [58], which was developed specifically for logic programming in Prolog. We note that the Mycroft and O’Keefe type system is considerably more complex, for example, formulas are annotated with types in addition to variables and constants.

We permit variables and constants to be annotated with construct-type and construct identifiers. That is, the set of permissible types for a source or target formula in a mapping rule consists of the ULD construct types and the constructs defined in the source or target configuration, respectively. Intuitively, if a variable in a mapping rule ranges over constructs, the variable is annotated with the appropriate construct type. Similarly, if a variable in a mapping rule ranges over construct instances, the variable is annotated with the appropriate construct. (Note that as an extension, we could also consider annotating construct instances with other construct instances according to conformance and *d-inst* relationships.)

$$\begin{array}{c}
(1) \frac{c/\text{ct-inst}(p, s)}{c/p :: s} \\
(2) \frac{c/\text{c-inst}(x, p)}{c/x :: p} \\
(3) \frac{c/\text{c-inst}(x, p), c \models \text{ct-inst}(p^*, s)}{c/p :: s} \\
(4) \frac{c/x.s:y, c/x :: p, c \models p^*.s^* \rightarrow q}{c/y :: q} \\
(5) \frac{c/\text{d-inst}(x, y), c/x :: p, c \models \text{conf}(p^*, q, d, r)}{c/y :: q} \\
(6) \frac{c/y :: q, c \models \text{unionof}(p, q^*)}{c/y :: p}
\end{array}$$

Figure 4.5: The basic type annotation rules for mappings.

$$\begin{array}{c}
(7) \frac{c/y \in x, c/x :: p, c \models \text{setof}(p^*, q)}{c/y :: q} \\
(8) \frac{c/y \in x, c/x :: p, c \models \text{listof}(p^*, q)}{c/y :: q} \\
(9) \frac{c/y \in x, c/x :: p, c \models \text{bagof}(p^*, q)}{c/y :: q} \\
(10) \frac{c/x[i] = y, c/x :: p, c \models \text{listof}(p^*, q)}{c/y :: q} \\
(11) \frac{c/y \in x=l, c/x :: p, c \models \text{bagof}(p^*, q)}{c/y :: q}
\end{array}$$

Figure 4.6: The collection type annotation rules for mappings.

|                                                                  |                                                                  |
|------------------------------------------------------------------|------------------------------------------------------------------|
| (12) $\frac{c/x[i] = y}{c/i :: \text{uldInteger}}$               | (13) $\frac{c/ x  = l}{c/l :: \text{uldInteger}}$                |
| (14) $\frac{c/y \in x = l}{c/l :: \text{uldInteger}}$            | (15) $\frac{c/x.s:y}{c/s :: \text{uldString}}$                   |
| (16) $\frac{c/p.s \rightarrow q}{c/s :: \text{uldString}}$       | (17) $\frac{c/\text{conf}(p, q, d, r)}{c/d :: \text{uldString}}$ |
| (18) $\frac{c/\text{conf}(p, q, d, r)}{c/r :: \text{uldString}}$ | (19) $\frac{\mathcal{K}_l^{v_1, \dots, v_n}(v)^t}{v :: t}$       |

Figure 4.7: The selector, conformance-cardinality, and identifier-generation type annotation rules for mappings.

The first type annotation rule states that, if  $p$  (a constant or variable) is a *ct-inst* of  $s$  (expressed as a formula in a mapping rule), the type of  $p$  is  $s$ . The second rule states that if  $x$  is a *c-inst* of  $p$ , the type of  $x$  is  $p$ . The third rule accesses data-model information for the case when a *ct-inst* relation is not explicitly given in the mapping rule. To illustrate, the following annotations result from applying the second and third rules to the previous mapping (Example 11).

$er/A :: \text{attribute}$   
 $er/\text{attribute} :: \text{struct-ct}$   
 $er/\text{uldString} :: \text{atomic-ct}$   
 $rel/T :: \text{relation}$   
 $rel/N :: \text{uldString}$   
 $rel/\text{relation} :: \text{struct-ct}$   
 $rel/\text{uldString} :: \text{atomic-ct}$

The fourth typing rule follows from axiom 27 (of Chapter 2) and states that the type of a component-selector value is the type of the corresponding component-definition value. To apply this annotation rule to the previous mapping, we must obtain the component definition in the corresponding source configuration. Applying the rule to the mapping in Example 11 adds the type annotation  $er/N :: \text{uldString}$ .

The fifth annotation rule follows from axiom 2 and 22 (as discussed in Chapter 2). Namely, we can use the conformance relationship in a configuration to determine the type of the construct instance serving as schema in a *d-inst* relationship. The sixth annotation rule enumerates union types. Rules seven through eleven assign collection types to members. Rules twelve through fourteen associate the index of a list, length of a collection, and number of repeated bag elements, respectively, to an integer atomic type. Rules fifteen and sixteen assign an atomic type to selectors; and rules seventeen and eighteen assign atomic types to conformance cardinality constraints. Finally, rule nineteen associates an atomic type (such as *uldURI*) to a generated identifier, if such an atomic type is given. Note that we do not specify a particular configuration for the generated identifier's type. Instead, we treat the type annotation as being applicable in any configuration in which the variable is used.

Once type annotations for a mapping rule are determined, we use the typing rules in Figure 4.8 to check for well-typed mappings. The typing rules are expressed using first-order logic and differ from the annotation rules in that they define constraints that must be satisfied for the mapping to be well-typed. The typing rules use notation similar to annotation rules, namely, the expression  $c/b$  is matched against formulas in the mapping, and the expression  $c \models b$  is matched against data-model information.

The first five typing rules require component selectors, *d-inst* expressions, membership expressions, list index expressions, and bag membership expressions, respectively, to be well-typed. The sixth rule requires every variable or constant with multiple types to have at least one non-union type. The seventh rule requires at most one such non-union type for each variable or constant. The eighth rule requires each type for variables or constants with more than one type to be related through a union definition. Note that the *unionConnected* constraint is defined by axiom 57 of Chapter 2. The ninth rule states that every variable and constant in a mapping rule (with a type annotation) has a ground type (that is, the type is either a variable with a known assignment or a constant). We assume the *ground* function returns true if its argument is ground, and false otherwise. Finally, the last typing rule restricts variables and constants with more than one distinct *atomic-ct* type (regardless of the configuration) to be type compatible. For example,

- (i)  $\forall(c, x, s, y) (c/x.s:y) \rightarrow \exists(p, q) (c/x :: p) \wedge (c/y :: q) \wedge (c/s :: \text{uldString}) \wedge (c \models p.s \rightarrow q)$
- (ii)  $\forall(c, x, y) (c/d\text{-inst}(x, y)) \rightarrow \exists(p, q, d, r) (c/x :: p) \wedge (c \models \text{conf}(p, q, d, r)) \wedge (c/y :: q) \wedge (c/d :: \text{uldString}) \wedge (c/r :: \text{uldString})$
- (iii)  $\forall(c, x, y) (c/y \in x) \rightarrow \exists(p, q) (c/x :: p) \wedge [c \models (\text{setof}(p, q) \vee \text{listof}(p, q) \vee \text{bagof}(p, q))] \wedge (c/y :: q)$
- (iv)  $\forall(c, x, y, i) (c/x[i]=y) \rightarrow \exists(p, q) (c/x :: p) \wedge (c \models \text{listof}(p, q)) \wedge (c/y :: q) \wedge (c/i :: \text{uldInteger})$
- (v)  $\forall(c, x, y, l) (c/y \in x=l) \rightarrow \exists(p, q) (c/x :: p) \wedge (c \models \text{bagof}(p, q)) \wedge (c/y :: q) \wedge (c/l :: \text{uldInteger})$
- (vi)  $\forall(c, v, t_1, \dots, t_n) (c/v :: t_1) \wedge (c/v :: t_2) \wedge \dots (c/v :: t_n) \rightarrow \exists(i) (1 \leq i \leq n) \wedge (c \models \neg \text{ct-inst}(t_i, \text{union-ct}))$
- (vii)  $\forall(c, v, t_1, t_2) (c/v :: t_1) \wedge (c/v :: t_2) \wedge (c \models \neg \text{ct-inst}(t_1, \text{union-ct})) \wedge (c \models \neg \text{ct-inst}(t_2, \text{union-ct})) \rightarrow (t_1=t_2)$
- (viii)  $\forall(c, v, t_1, t_2) (c/v :: t_1) \wedge (c/v :: t_2) \wedge (t_1 \neq t_2) \rightarrow \text{unionConnected}(t_1, t_2)$
- (ix)  $\forall(c, v, t) (c/v :: t) \rightarrow \text{ground}(t)$
- (x)  $\forall(c_1, c_2, v, t_1, t_2) (c_1/v :: t_1) \wedge (c_2/v :: t_2) \wedge (c_1 \models \text{ct-inst}(t_1, \text{atomic-ct})) \wedge (c_2 \models \text{ct-inst}(t_2, \text{atomic-ct})) \rightarrow \text{atomicCompatible}(t_1, t_2)$

Figure 4.8: The typing rules for ULD mappings.

**pcdata** and **uldString** are compatible atomic types whereas **uldInteger** and **uldString** are not. Note that, in the ULD, we do not provide an explicit mechanism for defining atomic-type compatibility, however, we assume that compatibility can be computed, for example, based on an external description.

To illustrate the use of typing rules, consider the following mapping rule that attempts to map E-R entities to relations in the relational data model. Using the type annotation rules, the variable AS is annotated with the construct **attSet**, which is defined as a *set-ct* construct in the E-R data model. However, the formula AS.hasName:N in the body of the rule treats AS as though it were a *struct-ct* construct. This formula violates the first typing rule in Figure 4.8, which states that a formula  $x.s:y$  is valid if the type of  $x$  has a component-selector definition for  $s$ . Here, since **attSet** is a set construct, the type for AS cannot have a component-selector definition.

$\text{c-inst}(\text{T}, \text{relation}), \text{T.hasName:N} \Leftarrow$

c-inst(E, entity), d-inst(E, T), T.hasAtts:AS, AS.hasName:N.

Similarly, the following mapping rule is also not well-typed. Unlike the previous example, the following mapping may result in an inconsistent target configuration, because the variable  $T$  in the head of the mapping rule is annotated with the two types, relation and attList, neither of which are union constructs (violating the seventh typing rule in Figure 4.8.)

$$\begin{aligned} &\text{c-inst}(T, \text{relation}), T.\text{hasAtts:AS}, \text{c-inst}(T, \text{attList}) \Leftarrow \\ &\text{c-inst}(E, \text{entity}), \text{d-inst}(E, T), T.\text{hasAtts:AS}. \end{aligned}$$

Finally, while not discussed here, the annotation and typing rules can also be applied directly to mapping rules that use an intensional formula. For mapping rules that use intensional formulas, we can determine type annotations by recursively applying the annotation rules to the variables and constants of the Datalog rules that define the intensional predicate. In other words, the type annotation rules can be directly applied to the Datalog rules that construct intensional predicates in a transformation.

### 4.3.2 Mapping Rules and Data-Model Constraints

Constraint macros offer a class of data-model constraints that can be used directly to determine whether a mapping rule is satisfiable with respect to the available constraint macros. Because there are a finite number of constraint macros, we can define specific procedures for each macro to check whether a mapping violates a target configuration.

Here we consider configuration-based uniqueness constraints, which take the form  $(p_1, p_2, \dots, p_n \text{ keyOf } c)$ , where  $n \geq 1$  and  $c$  is a *struct-ct* construct (or possibly a *union-ct* construct). Given a mapping rule  $h \Leftarrow b$ , where  $h$  and  $b$  each consist of one or more formulas, executing the rule may result in an inconsistent target configuration if each of the following conditions hold.

1. A set of component-selector formulas are present in the head of the rule, where each target formula originates from the same variable or constant  $x$  in configuration  $f_t$ , and the selectors represent a configuration-based uniqueness constraint. That is,

$\{f_t/x.p_{t1}:v_{t1}, f_t/x.p_{t2}:v_{t2}, \dots, f_t/x.p_{tn}:v_{tn}\} \subseteq h$ ,  $f_t/x :: c_t$ , and  $(p_{t1}, p_{t2}, \dots, p_{tn} \text{ keyOf } c_t)$  is a constraint in  $f_t$ .

2. A set of component-selector formulas are present in the body of the rule, where each formula originates from the same variable or constant  $y$  in configuration  $f_s$ , such that  $y$  is mapped to  $x$  and each selector value of  $y$  is mapped to a corresponding selector value for  $x$ . That is,  $\{f_s/y.p_{s1}:v_{s1}, f_s/y.p_{s2}:v_{s2}, \dots, f_s/y.p_{sn}:v_{sn}\} \subseteq b$ , and  $y=x$  or  $\mathcal{K}_l^y(x)$ ,  $v_{s1}=v_{t1}$  or  $\mathcal{K}_{l_1}^{v_{s1}}(v_{t1})$ ,  $v_{s2}=v_{t2}$  or  $\mathcal{K}_{l_2}^{v_{s2}}(v_{t2})$ , ..., and  $v_{sn}=v_{tn}$  or  $\mathcal{K}_{l_n}^{v_{sn}}(v_{tn})$ .
3. The type of  $y$  in (2) does not have a corresponding configuration-based uniqueness constraint. That is,  $f_s/y :: c_s$ , and  $(p_1, p_2, \dots, p_m \text{ keyOf } c_s)$  is not a constraint in  $f_s$  such that  $\{p_{s1}, p_{s2}, \dots, p_{sn}\} \subseteq \{p_1, p_2, \dots, p_m\}$ .

To illustrate the use of these conditions, consider the following mapping rule (taken from Example 11).

$$\begin{aligned} &\text{c-inst}(T, \text{relation}), T.\text{hasName:N}, \text{c-inst}(N, \text{uldString}) \Leftarrow \\ &\text{c-inst}(A, \text{attribute}), A.\text{hasName:N}, \mathcal{K}^A(T). \end{aligned}$$

As previously mentioned, the mapping rule may result in an inconsistent target configuration because relation names are required to be unique, whereas attribute names are not. The following shows that each of the previous conditions hold for the rule, and thus, the mapping may result in an inconsistent target configuration. (Note that the rule results in an inconsistent target configuration only if there is more than one attribute in the source.)

1.  $\{rel/T.\text{hasName:N}\} \subseteq h$ ,  $rel/T :: \text{relation}$ , and  $(\text{hasName keyOf relation})$  is a constraint in  $rel$ .
2.  $\{er/A.\text{hasName:N}\} \subseteq b$ ,  $\mathcal{K}^A(T)$ , and  $N=N$ .
3.  $rel/A :: \text{attribute}$  and  $(\text{hasName keyOf attribute})$  is not a constraint in  $er$ .

We note that checking constraint satisfiability for mapping rules requires (1) propagating constraints surrounding formulas in the body of the rule to formulas in the head of the

rule, and (2) ensuring that the propagated constraints do not violate the constraints associated with the target data models. For arbitrary ULD data-model constraints expressed using stratified Datalog $\neg$ , checking constraint satisfiability requires checking query satisfiability [50, 67] (that is, checking that the target query, the propagated constraints, and the existing constraints are satisfiable), which in general is undecidable [50, 51]. However, a number of useful constraints can still be exploited, such as the uniqueness constraint shown here.

### 4.3.3 Cardinality Constraints on Conformance

Similar to constraint macros, we can use cardinality constraints on conformance relationships to determine whether a mapping rule can generate an inconsistent target configuration. Here, we briefly sketch the approach. Consider the following mapping rule (changed slightly from `er-rel-8` to make it inconsistent).

$$\begin{aligned} &\text{c-inst}(P, \text{tuple}), \text{c-inst}(B, \text{table}), P \in B, \text{d-inst}(P, R) \Leftarrow \\ &\text{c-inst}(E, \text{entity}), \text{d-inst}(E, T), \mathcal{K}_{\text{tuple}}^E(P), \mathcal{K}_{\text{table}}^T(B), \mathcal{K}_{\text{relation}}^T(R). \end{aligned}$$

As shown, the rule generates a relational tuple for every entity in an E-R configuration. This mapping rule can result in an inconsistent target configuration. In particular, the E-R data model permits entities to be instances of multiple entity types, whereas the relational model requires a tuple to be an instance of exactly one type.

The following type annotations are generated for the mapping rule.

```

er/E :: entity
er/T :: entityType
rel/P :: tuple
rel/R :: relation
rel/B :: table

```

In addition, the formula  $\text{d-inst}(E, T)$  in the body of the rule and the formula  $\text{d-inst}(P, R)$  in the head of the rule assert that  $E$  is a data instance of  $T$ , and  $P$  is a data instance of  $R$ . We also know from the formulas  $\mathcal{K}_{\text{tuple}}^E(P)$  and  $\mathcal{K}_{\text{relation}}^T(R)$  that  $E$  generates the

identifier P, and T generates the identifier R. Therefore, the cardinality restriction on the conformance relationship between entities and entity types must not violate the cardinality restriction on the conformance relationship between tuples and relations.

The following definitions are taken from the data-model descriptions in Chapter 3 and state the conformance cardinality constraints for E-R entities and relational tuples. As shown, tuples must be data instances of exactly one relation, however, entities can be data instances of any number of entity types. Thus, the conformance constraint in the head of the mapping rule may be violated by the conformance constraints in the body of the rule.

```
construct tuple = list of uldValue conf(domain=*,range=1):relation
construct entity = {hasProps->propSet} conf(domain=*,range=*):entityType
```

In this case, the cardinality restriction in the head of the rule forms a specific uniqueness constraint that may be violated by the body of the rule. The following mapping, however, is satisfiable with respect to cardinality constraints on conformance. This mapping is exactly *er-rel-8*.

$$\begin{aligned} &\text{c-inst}(P, \text{tuple}), \text{c-inst}(B, \text{table}), P \in B, \text{d-inst}(P, R) \Leftarrow \\ &\text{c-inst}(E, \text{entity}), \text{d-inst}(E, T), \mathcal{K}_{\text{tuple}}^{E,T}(P), \mathcal{K}_{\text{table}}^T(B), \mathcal{K}_{\text{relation}}^T(R). \end{aligned}$$

#### 4.3.4 Inference and Transformation

We end this chapter by giving a simple example of automatically filling-in partial mapping rules using ULD axioms. We then briefly outline how a general approach might proceed when additional constraints are considered.

The following (Example 12) is a simple example of a partial mapping rule. In particular, the mapping rule does not specify what the variable N is an instance of. The rule is well-typed, but can result in an inconsistent target configuration because every identifier (in this case N), other than a construct type, must be defined as a construct or a construct instance. In this case, the formula  $\text{c-inst}(N, \text{uldString})$  should be added to the head of the rule.

**Example 12** (*A partial mapping rule*)

$$c\text{-inst}(\text{R}, \text{relation}), \text{R.hasName:N} \Leftarrow c\text{-inst}(\text{T}, \text{entityType}), \text{T.hasName:N}.$$

We use the following *elaboration* rules to fill-in *c-inst* relationships in partial mappings. The rules are similar to the type annotation rules of Figure 4.8 in that they are applied directly to mapping rules. However, instead of annotating constants and variables with types, an elaboration rule adds a *c-inst* formula to the head of a mapping rule. An elaboration rule takes the form **if**  $h$  **then**  $c/c\text{-inst}(x,p)$ , where  $h$  is a list of formulas and  $c$  is a target configuration. If the set of formulas  $h$  can be unified with the head of a mapping rule and data-model information from a configuration (rule eleven is a special case that we describe below), the formula  $c/c\text{-inst}(x,p)$  is added to the head of the mapping rule if it is not already present. Note that the elaboration rules below rely on type annotations and data-model information.

- (1). **if**  $c/x.s:y, c/x :: p, c \models p^*.s^* \rightarrow q$  **then**  $c/c\text{-inst}(y, q)$
- (2). **if**  $c/d\text{-inst}(x,y), c/x :: p, c \models \text{conf}(p^*, q, d, r)$  **then**  $c/c\text{-inst}(y, q)$
- (3). **if**  $c/y :: q, c \models \text{unionof}(p, q^*)$  **then**  $c/c\text{-inst}(y, p)$
- (4). **if**  $c/y \in x, c/x :: p, c \models \text{setof}(p^*, q)$  **then**  $c/c\text{-inst}(y, q)$
- (5). **if**  $c/y \in x, c/x :: p, c \models \text{listof}(p^*, q)$  **then**  $c/c\text{-inst}(y, q)$
- (6). **if**  $c/y \in x, c/x :: p, c \models \text{bagof}(p^*, q)$  **then**  $c/c\text{-inst}(y, q)$
- (7). **if**  $c/x[i] = y, c/x :: p, c \models \text{listof}(p^*, q)$  **then**  $c/c\text{-inst}(y, q)$
- (8). **if**  $c/y \in x=l, c/x :: p, c \models \text{bagof}(p^*, q)$  **then**  $c/c\text{-inst}(y, q)$
- (9). **if**  $c/x[i] = y$  **then**  $c/c\text{-inst}(i, \text{uldInteger})$
- (10). **if**  $c/y \in x=l$  **then**  $c/c\text{-inst}(l, \text{uldInteger})$
- (11). **if**  $\mathcal{K}_l^{v_1, \dots, v_n}(v)^t$  **then**  $c/c\text{-inst}(v, t)$

Each elaboration rule is applied directly to formulas in the head of the mapping except for the last rule (which is a special case). In particular, the last rule is matched against formulas in the body of the mapping and is only applied if  $v$  is used in a formula in the head of the mapping.

To illustrate, each formula in the consequent of the first elaboration rule holds for the mapping in Example 12. That is,  $T.\text{hasName}:N$  (from the rule body),  $T::\text{entityType}$  (from the annotation rule), and  $\text{entityType}.\text{hasName} \rightarrow \text{uldString}$  (from the configuration) are each true. Thus, the formula in the consequent of the elaboration-rule,  $c\text{-inst}(N, \text{uldString})$ , is added to the head of the mapping.

Finally, we consider the more general case of filling in partial mapping rules (as opposed to filling in only missing *c-inst* relationships, as above). Here, the goal is to allow a user to specify an initial correspondence assertion [64]—for example, that E-R entity types should be mapped to relations in the relational model—and based on the correspondence enumerate the remaining details of the mapping. In general, it is not possible to automatically define the desired mapping: (1) at some point, component selector labels must be unified, and (2) data-model constraints must be used to determine the validity of and help fill-in mappings, which is not a decidable problem for all constraints expressed in stratified Datalog $\neg$  (as discussed above). Instead of automatically generating the filled-in mapping, we sketch a procedure to fill in many, but potentially not all, details of the mapping.

The following is a list of approaches that can elaborate an initial mapping rule. Note that each approach below would be repeatedly applied to each formula in a single mapping rule  $h \Leftarrow b$  (until a fixed-point is reached), where both  $h$  and  $b$  consist of one more formulas. For example, given the statement “E-R entity types should be mapped to relations in the relational model,” we would construct the initial mapping rule  $c\text{-inst}(E, \text{relation}) \Leftarrow c\text{-inst}(E, \text{entityType})$ , and apply the following techniques to fill the rule in.

**Structural Elaboration** Given a formula  $c/c\text{-inst}(x, p)$  in either  $h$  or  $b$  (of the mapping rule) where  $ct\text{-inst}(p, \text{struct-ct})$  is in configuration  $c$ , we can add the formulas  $c/x.s:t$  and  $c/c\text{-inst}(y, q)$  to  $h$  or  $b$  (where  $s$  and  $q$  are constants and  $t$  is a variable), respectively, whenever  $y.s \rightarrow q$  is in  $c$ . We call this approach *structural elaboration*, because each component-selector definition must have a corresponding value in a construct instance.

**Conformance Elaboration** Given a formula  $c/c\text{-inst}(x, p)$  in either  $h$  or  $b$  where  $\text{conf}(p,$

$q, d, r$ ) is in  $c$  and  $r$  is either ‘1’ (exactly one) or ‘+’ (one or more), we add the formulas  $c/d\text{-inst}(x, y)$  and  $c/c\text{-inst}(y, q)$  to  $h$  or  $b$ , respectively (where  $q$  is a constant and  $y$  is a variable). Similarly, given a formula  $c/c\text{-inst}(y, q)$  in either  $h$  or  $b$  where  $\text{conf}(p, q, d, r)$  is in  $c$  and  $d$  is either ‘1’ (exactly one) or ‘+’ (one or more), we add the formulas  $c/d\text{-inst}(x, y)$  and  $c/c\text{-inst}(x, p)$  to  $t$  or  $b_m$ , respectively (where  $p$  is a constant and  $x$  is a variable). We call this approach *conformance elaboration*, because each construct instance with a required cardinality constraint on conformance requires an instance of (or a type for) at least one other construct instance.

**Constraint Elaboration** Given the constraint  $e_1, \dots, e_i \Leftarrow b_{c1}, \dots, b_{cj}$  in configuration  $c$ , we can add the formulas  $c/e_1, \dots, c/e_i$  to  $h$  if  $\{c/b_{c1}, \dots, c/b_{cj}\} \subseteq h$  with variable and constant substitution. In other words, if the body of the constraint matches (with unification) a subset of formulas in the head of the mapping rule, the constraint formulas  $e_1$  to  $e_i$  are added (with variable substitution) to the head of the mapping rule. This technique can be similarly applied to the body of the mapping rule. We call this approach *constraint elaboration*, because we add information to the mapping based on available constraints.

The following six steps illustrate how the elaboration approaches above can be used to fill-in the initial mapping rule  $c\text{-inst}(E, \text{relation}) \Leftarrow c\text{-inst}(E, \text{entityType})$ . Once the rule is fully elaborated (the result of step 6), the user must still state the connection between identifiers, for example, that the relation name and the entity name should be the same (N1 and N2 below).

**Step 1.** Apply the initial mapping rule.

$$c\text{-inst}(E, \text{relation}) \Leftarrow c\text{-inst}(E, \text{entityType}).$$

**Step 2.** Apply structural elaboration to the result of step 1 to obtain the following rule.

$$\begin{aligned} c\text{-inst}(E, \text{relation}), E.\text{hasName:N2}, E.\text{hasAtts:AL} \Leftarrow \\ c\text{-inst}(E, \text{entityType}), E.\text{hasName:N1}, E.\text{hasAtts:AS}. \end{aligned}$$

**Step 3.** Apply conformance elaboration to the result of step 2 to obtain the following rule.

$$\begin{aligned} & \text{c-inst}(E, \text{relation}), E.\text{hasName:N2}, E.\text{hasAtts:AL}, \text{d-inst}(B, E), \text{c-inst}(B, \text{table}) \Leftarrow \\ & \text{c-inst}(E, \text{entityType}), E.\text{hasName:N1}, E.\text{hasAtts:AS}. \end{aligned}$$

**Step 4.** Apply constraint elaboration to the result of step 3 to obtain the following rule. Here, we apply the E-R constraint (er-36):  $A \in AS \Leftarrow \text{c-inst}(E, \text{entityType}), E.\text{hasAtts:AS}$ .

$$\begin{aligned} & \text{c-inst}(E, \text{relation}), E.\text{hasName:N2}, E.\text{hasAtts:AL}, \text{d-inst}(B, E), \text{c-inst}(B, \text{table}) \Leftarrow \\ & \text{c-inst}(E, \text{entityType}), E.\text{hasName:N1}, E.\text{hasAtts:AS}, A1 \in AS. \end{aligned}$$

**Step 5.** Apply constraint elaboration to the result of step 4 to obtain the following rule. Here, we apply the relational constraint (rel-1): (attList isNonEmpty). Note that the constraint macro is expanded to an expression in the ULD constraint language prior to its application.

$$\begin{aligned} & \text{c-inst}(E, \text{relation}), E.\text{hasName:N2}, E.\text{hasAtts:AL}, \text{d-inst}(B, E), \text{c-inst}(B, \text{table}), \\ & A2 \in AL, \text{c-inst}(A2, \text{attribute}) \Leftarrow \\ & \text{c-inst}(E, \text{entityType}), E.\text{hasName:N1}, E.\text{hasAtts:AS}, A1 \in AS. \end{aligned}$$

**Step 6.** Apply structural elaboration to the result of step 5 to obtain the following rule, which is the fixed-point.

$$\begin{aligned} & \text{c-inst}(E, \text{relation}), E.\text{hasName:N2}, E.\text{hasAtts:AL}, \text{d-inst}(B, E), \text{c-inst}(B, \text{table}), \\ & A2 \in AL, \text{c-inst}(A2, \text{attribute}), A2.\text{hasName:N4}, A2.\text{hasDomain:D2} \Leftarrow \\ & \text{c-inst}(E, \text{entityType}), E.\text{hasName:N1}, E.\text{hasAtts:AS}, A1 \in AS, A1.\text{hasName:N3}, \\ & A1.\text{hasDomain:D1}. \end{aligned}$$

## 4.4 Summary

In this chapter, we described the ULD transformation language for converting information between data models, schemas, and instance data. We used the language to describe

standard data-model mappings, including E-R-to-relational and XML-to-relational transformations. We also used the transformation language to express a schema mapping, where the source and target configurations are represented in distinct data models. Finally, we outlined a number of approaches to ensure transformation mappings are valid, and described two simple approaches to elaborate partial transformation rules.

We have used the ULD transformation language to implement a number of mappings. We briefly describe our experience using the language here. In Chapter 6, we discuss ULD implementation issues, including transformation. We note that, in general, there are a number of opportunities for optimizing transformations, which we do not consider in this dissertation (see Section 6.1). However, based on our experience, we believe that transformation can be practically implemented using existing languages such as Prolog.

Our first experiments with transformation were done within the context of superimposed information models [36, 37]. Our goal was to enable interoperability for *model-based* tools. By *model-based*, we mean a tool that works for any valid information represented in a particular data model, as opposed to a tool that relies on a particular schema (in a particular data model). By defining and executing data-model mappings, our goal was to allow a user to select data of interest independently from the tool they wish to use. Thus, a user would be free to use their favorite tool over data represented in an arbitrary data model. We defined and executed a number of mappings (using Prolog) for converting information represented in the SLIMPad data model [16, 37] to information represented in the RDF and Topic-Map data models. The SLIMPad data model uses simple tree-based structures for storing data, similar to XML. We note that in these mappings, source configurations were small in size, the mapping rules did not require recursion or negation, and the transformations did not define one-to-one data-model conversions.

We also defined and executed schema mappings, where the source and target schemas were represented using different data models (XML and Topic Maps, respectively). These experiments used larger data sets, where each source configuration contained between eight- and twelve-thousand ULD facts. Each XML source contained metadata for a particular section of an on-line course, such as the section's slides (where each slide was considered an addressable, external resource), audio-video files, homework assignments,

reading assignments, and so on. The XML files were used to build an HTML version of an on-line course. However, the XML files were not convenient for course developers who wished to reuse slides and other materials when developing new courses or updating existing courses. The goal of the target Topic-Map schema was to serve as a simple, conceptual model that course developers could navigate (using a Topic-Map browser) to easily find material of interest.

We originally implemented these mappings using Prolog, however, as the size of the source configurations increased, the system became slow for mapping rules with recursion and negation. (Note that only a few mapping rules used recursion.) Based on these results, we developed a prototype system, written in Java, that stored configurations in a relational database as opposed to a file of Prolog facts. We defined a relational schema to represent the ULD facts and converted mapping rules to embedded SQL programs. (Mapping rules with recursion and negation were converted to embedded SQL by hand.) Using SQL and storing configurations in a relational database improved the overall speed of executing the transformation rules.

Finally, we also used the ULD transformation language to execute schema mappings where the source and target configurations were both expressed using XML. The two schemas described similar intrusion-detection information, for example, that would be generated by an intrusion-detection system, and served as interchange formats for use in other security applications. The schemas contained various forms of structural heterogeneity (at the schema level), including different representations for data as well as structural and schematic heterogeneity. Using the ULD transformation language, we were able to declaratively describe the schema mapping, and using Prolog, execute the mapping directly to convert data between the schemas.

# Chapter 5

## Incremental Navigation

Like transformation, navigation is an application of the ULD. The goal of navigation is to provide generic access to structured information, much like a Web browser provides generic access to viewable information. We are particularly interested in browsing a data source where type information (if available) and structural relationships among the data items can be viewed in a step-by-step manner, where a user can select an item, select a path that leads from the item, follow the path to a new item, and so on, *incrementally* through the source. The need for incremental navigation is motivated by the following uses. First, we believe that simple browsing tools provide people with a powerful and easy way to become familiar with an information source. Additionally, having generic access to heterogeneous sources supports exploration tools used in the process of integration [25, 34]. In particular, once an information source has been identified, its contents can be examined (by a person or an agent) to determine if and how it should be combined (or integrated) with other sources.

In this section, we describe a generic set of incremental navigation operators that are implemented against the ULD, which enables various styles of browsing across a wide range of data models. In addition, the ULD provides multiple ways to specify how navigation operators should work with various data models. We consider both a low-level approach for creating detailed and complete specifications as well as a simple, high-level approach for semi-automating each operator's implementation.

## 5.1 Motivating Examples

Incremental navigation can help an information agent (such as a crawler) or a person to become familiar with a new data source. When an information agent discovers a new source (that uses the schema, for example, in Figure 1.3), it may wish to know: (1) what data model is used (that is, is it an RDF, XML, Topic Map, or relational source?), (2) assuming RDF is used, whether any classes are defined for the source (what is the source schema?), (3) which properties are defined for a given class (what properties does the *film* class have?), (4) which objects exist for the class (what are the instances of the *film* class?) and (5) what kinds of values exist for a given property of a particular object of the class (what *actor* objects are involved in this *film* object?). The agent moves step-by-step through the source to gather information. It starts by determining what kind of source it is, and moves down through the schema to individual data items, following their links to other items.

This example assumes the user (or agent) understands the data model of the source. For example, the agent begins navigation by asking the source which RDF classes it stores. If the data model used was XML instead of RDF, the agent (assuming it understands the XML model) could have started navigation by asking for all of the available element types. We call this approach *data-model-aware browsing*, in which the constructs of the data model can be used to guide navigation.

We primarily focus in this chapter on providing incremental navigation to users who may not understand the underlying representation scheme (although we do consider data-model-aware browsing as well). We call this *direct browsing*, because the data and schema items of the source are navigated directly. Using direct browsing, our previous example would change as follows. Upon discovering an information source, the crawler might wish to know: (1) what kind of information the source contains, which (in our example) would include “films,” “actors,” and “awards,” etc., (2) (assuming the crawler is interested in films) the things that describe films, which would include “titles” and relationships to awards and actors, (3) the available films in the source, and (4) the actors of a particular film, which is obtained by stepping across the “involved” link for the film in question.

Thus, without requiring any knowledge of the source’s data model, with direct navigation, the agent can still freely browse the source to gather information.

## 5.2 Navigation Operators

Incremental navigation refers to the ability to step iteratively through an information source one item at a time. For example, when situated at a particular item (a *location*), we can use incremental operators to determine the set of available links from the current position and follow one of these links to a new location. This style of navigation is in contrast to approaches that navigate information using path expressions [29, 45], which describe collections of paths and variable bindings over a database. Path expressions enhance query capability, whereas our goal is to enable step-by-step navigation through an information source. The following issues motivate our choice of operators.

- *How should items be represented?* We wish to support navigation regardless of the data model used by the source. However, sources represent information items using different data-model constructs. At issue is whether to expose atomic (for example, text strings) or structured values (for example, records or lists) through navigation operations. We define the navigation operators such that each operator takes and returns atomic values, which keeps the navigation interface simple, and does not require additional knowledge of the underlying structures of data models.
- *How should navigation start?* By incremental navigation, we mean navigating from a current location to a new location. However, when a source is first encountered, it is not immediately clear from which item to start navigation. Path-based approaches typically use available collections or collections formed from queries to apply navigation expressions. Because the items for beginning navigation depend on the individual source and its data model, we choose to include an operator for obtaining the starting locations of sources (that is, the operator returns one or more *entry points*).

- *How should schema-instance relationships be handled?* Data models can have different requirements on schema-data conformance and these constraints have implications for navigation. For example, a partial-schema data model may require support for navigating instances that are not included in the schema, because the schema alone cannot be used to navigate all available data. We define a minimal set of generic operators that can be used to access the various styles of schema arrangements possible in data models.
- *What information within a source should be navigable?* For a given source, some information may be meaningful for navigation, while other information may not. For example, data-model-aware navigation might offer navigable items that are not necessarily useful in direct navigation. We allow different implementations for each navigation operator through navigation bindings, expressed using the ULD query language.

The incremental navigation interface consists of the following operators. Navigable items in an information source are separated into disjoint sets of *links* and *locations*. For a given configuration, we assume a finite set of location names  $\mathcal{L}$  and a finite set of *link* names  $\mathcal{N}$  where both  $\mathcal{L}$  and  $\mathcal{N}$  consist of simple, atomic string values.

- **Starting Points.** The operator  $sloc : \mathcal{P}(\mathcal{L})$  returns all available entry points into an information source. Note that we require the result of  $sloc$  to be a set of locations (as opposed to links).
- **Links.** The operator  $links : \mathcal{L} \rightarrow \mathcal{P}(\mathcal{N})$  returns all out-bound links available from a particular location. We note that for some locations, there may not be any links available, that is, the  $links$  operator may return the empty set.
- **Following Links.** The operator  $follow : \mathcal{L} \times \mathcal{N} \rightarrow \mathcal{P}(\mathcal{L})$  returns the set of locations that are at the end of a given link. We use the  $follow$  operator to prepare to move to a new location from our current location. Note that following a link may lead to a set of locations. Given the set of locations returned by the  $follow$  operator, the user or agent directing the navigation can choose one as the new location.

- **Types.** The operator  $types : \mathcal{L} \rightarrow \mathcal{P}(\mathcal{L})$  returns the (possibly empty) set of types for a given location. We use the  $types$  operator to obtain locations that represent the schema for a data item. Note that a particular location may have zero, one, or many associated types.
- **Extents.** The operator  $extent : \mathcal{L} \rightarrow \mathcal{P}(\mathcal{L})$  returns the (possibly empty) set of instances for a given location. The  $extent$  operator computes the inverse of the  $types$  operator.

We call a group of navigation-operator implementations (that is, one implementation for each operator) a *navigation binding* (or *binding* when the context is clear). To issue the  $sloc$  operation with binding  $b$  on configuration  $c$ , we use the notation  $b[c].sloc$ . To illustrate, consider the following binding  $B_{rdf}$  for navigating RDF sources (that is, this binding returns data-model information for data-model-aware browsing). The following shows a sequence of calls to  $B_{rdf}$ 's operator implementations and the results for a configuration named  $c$ .

$$\begin{aligned}
 B_{rdf}[c].sloc &= \{\text{'class'}, \text{'property'}, \text{'subClass'}, \text{'resource'}, \text{'triple'}, \dots\} \\
 B_{rdf}[c].extent(\text{'class'}) &= \{\text{'film'}, \text{'comedy'}, \text{'thriller'}, \dots\} \\
 B_{rdf}[c].links(\text{'thriller'}) &= \{\text{'title'}\} \\
 B_{rdf}[c].extent(\text{'thriller'}) &= \{\text{'#m1'}, \dots\} \\
 B_{rdf}[c].follow(\text{'#m1'}, \text{'title'}) &= \{\text{'The Usual Suspects'}\}
 \end{aligned}$$

The goal of the operators is to allow navigation through an information source incrementally, as opposed to issuing special-purpose queries to obtain the same information. Note that, as shown by the *follow* operator in the example above, bindings present a view of the data that does not require the user to connect the various underlying data structures (for example, by joining various structures together)—in this case, the *follow* operator binding computes the connection between a property name 'title' and the use of the property within a triple (to obtain a new location). The challenge of providing incremental navigation is in specifying bindings of general use for navigation, and simplifying the effort on the part of the user for stepping through a source.

### 5.3 Navigation Bindings as ULD Queries

We present two ways that bindings can be implemented against the ULD. The first method treats a binding as a set of queries using the ULD query language. We call these *binding queries*. The second method provides for higher-level description of bindings (and is described in the next section).

For the query approach, we treat an operator binding as the result of a ULD query where each operator is represented as an intensional predicate in a Datalog program. We represent the incremental navigation operators using the following intensional predicates.

- $\text{sloc}(r)$ , where  $r$  represents a starting location.
- $\text{links}(l, r)$ , where  $r$  is a link from location  $l$ .
- $\text{follow}(l, p, r)$ , where  $r$  is the location that is found by following link  $p$  from some location  $l$ .
- $\text{types}(l, r)$ , where  $l$  is a location of type  $r$ .
- $\text{extent}(l, r)$ , where  $r$  is in the extent of  $l$ .

For a given operator invocation, the binding computes the least fix-point of the operator's predicate over the configuration, and returns the associated set of results. The following set of rules define a portion of the binding  $B_{rdf}$  above. We note that the definition uses an additional intensional predicate for computing the transitive closure of the RDF subclass relationship.

$$\begin{aligned}
 \text{sloc}(R) &\leftarrow \text{ct-inst}(R, C). \\
 \text{extent}(\text{class}, R) &\leftarrow \text{c-inst}(R, \text{class}). \\
 \text{links}(L, R) &\leftarrow \text{c-inst}(C, \text{class}), C.\text{hasLabel}:L, \text{c-inst}(P, \text{property}), P.\text{hasLabel}:R, \\
 &\quad P.\text{hasDomain}:C', \text{subClassClosure}(C, C'). \\
 \text{extent}(L, R) &\leftarrow \text{c-inst}(C, \text{class}), C.\text{hasLabel}:L, \text{d-inst}(R, C). \\
 \text{follow}(L, P, R) &\leftarrow \text{c-inst}(T, \text{triple}), T.\text{hasPred}:X, T.\text{hasSubj}:L, T.\text{hasObj}:R, \\
 &\quad \text{c-inst}(R, \text{literal}), \text{d-inst}(L, C), \text{c-inst}(X, \text{property}), X.\text{hasLabel}:P, \\
 &\quad X.\text{hasDomain}:C', \text{subClassClosure}(C, C').
 \end{aligned}$$

$\text{subClassClosure}(C, C) \leftarrow \text{c-inst}(C, \text{class}).$   
 $\text{subClassClosure}(C1, C3) \leftarrow \text{c-inst}(S, \text{subClass}), S.\text{hasSub}:C1, S.\text{hasSuper}:C2,$   
 $\text{subClassClosure}(C2, C3).$

In addition to data-model-aware navigation, we can also define operator bindings for direct navigation using the same approach. For example, the following binding queries present a complete-schema view of an RDF configuration. As shown, the only starting locations are classes and the only available links from a class are its associated properties and its associated instances. (Note that we reuse the *subClassClosure* intensional predicate defined above.)

$\text{sloc}(R) \leftarrow \text{c-inst}(C, \text{class}), C.\text{hasLabel}:R.$   
 $\text{links}(L, R) \leftarrow \text{c-inst}(C, \text{class}), C.\text{hasLabel}:L, \text{c-inst}(P, \text{property}), P.\text{hasLabel}:R,$   
 $P.\text{hasDomain}:C', \text{subClassClosure}(C, C').$   
 $\text{links}(L, R) \leftarrow \text{c-inst}(C, \text{class}), \text{d-inst}(O, C), O.\text{hasURI}:L, \text{c-inst}(T, \text{triple}),$   
 $T.\text{hasPred}:X, T.\text{hasSubj}:O, \text{c-inst}(X, \text{property}), X.\text{hasLabel}:R,$   
 $X.\text{hasDomain}:C', \text{subClassClosure}(C, C').$   
 $\text{follow}(L, P, R) \leftarrow \text{c-inst}(C, \text{class}), C.\text{hasLabel}:L, \text{c-inst}(T, \text{property}), T.\text{hasLabel}:P,$   
 $T.\text{hasDomain}:C', T.\text{hasRange}:C'', C''.\text{hasLabel}:R,$   
 $\text{subClassClosure}(C, C').$   
 $\text{follow}(L, P, R) \leftarrow \text{c-inst}(C, \text{class}), C.\text{hasLabel}:L, \text{c-inst}(T, \text{property}), T.\text{hasLabel}:P,$   
 $T.\text{hasDomain}:C', T.\text{hasRange}:R, R = \text{'literal'},$   
 $\text{subClassClosure}(C, C').$   
 $\text{follow}(L, P, R) \leftarrow \text{c-inst}(C, \text{class}), \text{d-inst}(O, C), O.\text{hasURI}:L, \text{c-inst}(T, \text{triple}),$   
 $T.\text{hasPred}:X, T.\text{hasSubj}:O, T.\text{hasObj}:R, \text{c-inst}(R, \text{literal}),$   
 $\text{c-inst}(X, \text{property}), X.\text{hasLabel}:P, X.\text{hasDomain}:C',$   
 $\text{subClassClosure}(C, C').$   
 $\text{follow}(L, P, R) \leftarrow \text{c-inst}(C1, \text{class}), \text{d-inst}(O1, C1), O1.\text{hasURI}:L, \text{c-inst}(T, \text{triple}),$   
 $T.\text{hasPred}:X, T.\text{hasSubj}:O1, T.\text{hasObj}:O2, O2.\text{hasURI}:R,$   
 $\text{c-inst}(X, \text{property}), X.\text{hasLabel}:P, X.\text{hasDomain}:C',$   
 $\text{subClassClosure}(C, C').$   
 $\text{extent}(L, R) \leftarrow \text{c-inst}(C, \text{class}), C.\text{hasLabel}:L, \text{d-inst}(O, C), O.\text{hasURI}:R.$   
 $\text{type}(L, R) \leftarrow \text{c-inst}(C, \text{class}), C.\text{hasLabel}=R, \text{d-inst}(O, C), O.\text{hasURI}:L.$

```

RDF Binding  $B$  =  $(L, N, S, F)$ 
 $L$  = class, resource, literal
 $N$  = property, triple
 $S$  = class
 $F$  = triple : resource[triple.hasSubj]⇒literal[triple.hasObj],
      triple : resource[triple.hasSubj]⇒resource[triple.hasObj],
      property : class[property.hasDomain]⇒class[property.hasRange]
name(X,N) ← c-inst(X,class), X.hasLabel:N.
name(X,N) ← c-inst(X,resource), X.hasURI:N.
name(N,N) ← c-inst(N,literal).
name(X,N) ← c-inst(X,property), X.hasLabel:N.
name(X,N) ← c-inst(X,triple), X.hasPred=R, c-inst(R,property), R.hasLabel=N.

```

Figure 5.1: A high-level binding for direct navigation of RDF.

## 5.4 Higher-Level Binding Specifications

The binding-queries approach is a low-level mechanism for providing detailed and complete descriptions of bindings. Here we describe a higher-level mechanism to specify bindings, which we call the *binding constructs* approach. This approach is used for specifying direct-navigation bindings and allows a user to select certain constructs as *locations* (to be viewed) and certain other constructs as *links* (to be followed). Thus, a user can easily restrict the subset of links and locations as well as how they are traversed without having to write Datalog rules. Using this specification, the navigation operators are automatically computed by traversing the appropriate instances of locations and links in the configuration. For example, to compute the *links* operator from a given location, we start from the location's construct instance and consider connections from this instance to other instances to see if any are designated as a link. Figure 5.1 shows an example of a binding-constructs definition for RDF. In the example, RDF classes, resources, and literals are considered sources for locations and RDF properties and triples are considered sources for links.

We define the binding-constructs specification as a tuple of the form  $(L, N, S, F)$ . The disjoint sets  $L$  and  $N$  consist of construct identifiers such that  $L$  is the set of constructs used as locations and  $N$  is the set of constructs used as links. The set  $S \subseteq L$  gives the entry points of the binding. Finally, the set  $F$  contains *link definitions* (described below), which fill in the details of how links should be traversed.

Each construct in  $L$  and  $N$  has an associated naming definition that describes how to compute the name of an instance of the construct. The name would typically be viewed by the user during navigation. (Recall that the incremental operators work over string values, therefore, the naming definitions serve to map location and link instances to appropriate string values.) A naming definition is expressed as a Datalog rule whose head is the intensional predicate  $\text{name}(x,n)$  such that  $x$  is a given construct instance and  $n$  is the string to be used as the name of the instance. For example, in Figure 5.1, RDF classes and properties are named by their labels, resources are named by their URI values, a literal value is used directly as its name, and the name of a triple is the name of its associated predicate value.

The incremental operators in the binding-constructs approach are computed automatically by traversing connected instances. We define the following rules to compute when two instances are connected. (Note that these connected rules only perform single-step traversal, and in general, can be extended to allow an arbitrary number of steps.)

$$\begin{aligned} \text{connected}(X_1, X_2) &\leftarrow \text{c-inst}(X_1, C), \text{ct-inst}(C, \text{struct-ct}), X_1.S:X_2. \\ \text{connected}(X_1, X_2) &\leftarrow \text{c-inst}(X_1, C), \text{ct-inst}(C, \text{bag-ct}), X_2 \in X_1. \\ \text{connected}(X_1, X_2) &\leftarrow \text{c-inst}(X_1, C), \text{ct-inst}(C, \text{list-ct}), X_2 \in X_1. \\ \text{connected}(X_1, X_2) &\leftarrow \text{c-inst}(X_1, C), \text{ct-inst}(C, \text{set-ct}), X_2 \in X_1. \\ \text{connected}(X_2, X_1) &\leftarrow \text{connected}(X_1, X_2). \end{aligned}$$

A connected formula is true when there is a structural connection between two instances. Note that the rules above do not consider the case when two items are linked by a *d-inst* relationship. Instead, the *d-inst* relationship is directly used by the *types* and *extent* operators, whereas connections are used by the *links* and *follow* operators.

The link definitions of  $F$  serve to describe the behavior of links. A link definition in  $F$  has the form  $c_k : c_1[p_1] \Rightarrow c_2[p_2]$  where:

- The behavior of the link construct  $c_k \in N$  is being described by the rest of the expression. For example, in Figure 5.1, the first link definition is for triple constructs.
- For  $c_1, c_2 \in L$ , the construct  $c_k$  can serve to link an instance of  $c_1$  to an instance of  $c_2$ . Thus, we can traverse from instances of  $c_1$  to instances of  $c_2$  via an instance

of  $c_k$ . For example, in Figure 5.1, the first link definition says that we can follow a resource instance to a literal instance if they are connected by a triple.

- The expressions  $p_1$  and  $p_2$  further define how instances of  $c_k$  can be used to link  $c_1$  and  $c_2$  instances, respectively. For example, in Figure 5.1, the first link definition states that triples link resources and literals through the triple's `hasSubj` and `hasObj` selector, respectively.

Link definitions are used directly to define the *links* and *follow* operations. We introduce the  $linkSource(F, i_1, i_k)$  and  $linkTarget(F, i_k, i_2)$  clauses to filter possible connections that are considered when computing the *links* and *follow* operations. Given a set of link definitions  $F$  and a connection from  $i_1$  to  $i_k$  such that  $connected(i_1, i_k)$  is true (where  $i_k$  is the link instance),  $linkSource(F, i_1, i_k)$  is true if there is at least one link definition  $c_k : c_1[p_1] \Rightarrow c_2[p_2]$  in  $F$  such that  $i_1$  is an instance of  $c_1$  (that is,  $c-inst(i_1, c_1)$  is true),  $i_k$  is an instance of  $c_k$  (that is,  $c-inst(i_k, c_k)$  is true where  $c_k$  is a link construct), and  $i_1$  and  $i_k$  are connected according to the expression  $p_1$ . Similarly,  $linkTarget(F, i_k, i_2)$  is true if there is at least one link definition  $c_k : c_1[p_1] \Rightarrow c_2[p_2]$  in  $F$  such that  $i_k$  is an instance of  $c_k$  (a link construct),  $i_2$  is an instance of  $c_2$ , and  $i_k$  and  $i_2$  are connected according to the expression  $p_2$ .

Given the above definitions, we automatically compute each navigation operator using the Datalog queries in Figure 5.2. We assume each operator is represented as an intensional predicate (as before) and the binding specification  $B = (L, N, S, F)$  is stored as a set of extensional predicates  $G_l(X)$ ,  $G_n(X)$ ,  $G_s(X)$ , and  $G_f(X)$ . For example, the expression  $G_l(X)$  binds  $X$  to  $L$  in the binding  $B$ . We also assume that the *name* predicate is stored as an intensional formula (defined in the binding).

The first rule finds the set of entry points: It obtains a construct in the set of starting locations, finds an instance of the construct, and then computes the name of the instance. The second rule finds the locations with links. For each named instance in the configuration that is connected to another instance, we use the *linkSource* predicate to check if it is a valid connection, we check to make sure that the link instance (represented as the variable  $Y$ ) is valid, and then compute the name of the instance. The third rule is similar to the

$$\begin{aligned}
\text{sloc}(R) &\leftarrow G_s(S), P \in S, \text{c-inst}(X, P), \text{name}(X, R). \\
\text{links}(L, R) &\leftarrow \text{name}(X, L), \text{connected}(X, Y), G_f(F), \text{linkSource}(F, X, Y), \text{c-inst}(Y, P), \\
&\quad G_n(N), P \in N, \text{name}(P, R). \\
\text{follow}(L, P, R) &\leftarrow \text{name}(X, L), \text{connected}(X, Y), G_f(F), \text{linkSource}(F, X, Y), \text{c-inst}(Y, Q_1), \\
&\quad G_n(N), Q_1 \in N, \text{name}(Q_1, P), \text{connected}(Y, Z), \text{c-inst}(Z, Q_2), \\
&\quad G_l(L), Q_2 \in L, \text{linkTarget}(F, Y, Z). \\
\text{type}(L, R) &\leftarrow \text{name}(X, L), \text{c-inst}(X, P), G_l(L), P \in L, \text{d-inst}(X, Y), \text{c-inst}(Y, Q), Q \in L, \\
&\quad \text{name}(Y, R). \\
\text{extent}(L, R) &\leftarrow \text{name}(Y, L), \text{c-inst}(Y, Q), G_l(L), Q \in L, \text{d-inst}(X, Y), \text{c-inst}(X, P), P \in L, \\
&\quad \text{name}(X, R).
\end{aligned}$$

Figure 5.2: The Datalog rules to compute links and locations based on a binding-constructs specification.

second, except it additionally uses the *linkTarget* to determine the new location. Finally, the last two rules use the *d-inst* relationship to find types and extents.

To demonstrate the approach, we use the binding definition of Figure 5.1 and the sample configuration of Figure 5.3. This configuration shows part of Figure 1.3 as a graph whose nodes are construct instances and edges are either connections between structures or *d-inst* links. Consider the following series of invocations, assuming the configuration is named *c*.

1.  $B[c].\text{sloc} = \{\text{'film'}, \text{'thriller'}\}$ . According to the binding definition, the *sloc* operator returns all the labels of *class* construct instances. As shown in Figure 5.3, the only *class* construct instances are *thriller* and *film*.
2.  $B[c].\text{links}(\text{'film'}) = \{\text{'title'}\}$ . The *links* operator is computed by considering each connected construct instance of *film* until it finds a construct instance whose associated construct is in *N*. As shown, the only such instance is *title*, which is an RDF property.
3.  $B[c].\text{extent}(\text{'film'}) = \{\text{'#m1'}\}$ . The *extent* operator looks for the *d-inst* links of the given instance. As shown, the only such link for *film* is to *m1*.
4.  $B[c].\text{follow}(\text{'#m1'}, \text{'title'}) = \{\text{'The Usual Suspects'}\}$ . The *follow* operator starts in the same way as the *links* operator by finding instances (whose constructs are in *N*)




Figure 5.3: An example of labeled instances for RDF.

that are connected to the given instance. For the given item, the only such instance in Figure 5.3 is *t1*. The *follow* operator then returns the *hasObj* component of *t1*, according to the link definition for *triple*.

Note that according to the rules for computing the *follow* and *links* operators, we only consider instances that are directly connected to each other (through the *connected* predicate). Thus, invoking the operator *links*(‘thriller’) will not return a result (for our example) because there are no link construct instances directly connected to the associated RDF class. However, because of the implied relationship among subclasses, the algorithm should also include the *title* link, because *thriller* is a subclass of *film*. That is, the properties of a class in RDF also include the properties of its subclasses. One way to include such information is to expand the set of connection rules to consider subclass instances. The connection rules could be augmented to consider instances that are indirectly connected through intermediate instances (for example, for certain cases such as subclass). These new connection definitions could be added directly to the binding definition of the RDF data model. Alternatively, it may be possible to extend the binding specification to include multiple steps, for example, through arbitrary path expressions in *F*, which could be used to define the appropriate connected rules.

Using the ULD, we have built a prototype tool for data-model-aware browsing called

*JustBrowsing* [18]. In *JustBrowsing*, we implemented a graphical user interface and the navigation operations to allow users to browse information stored in the RDF, relational, XML, and Topic-Map data models (see Section 6.1). Although not discussed here, we believe that by allowing multiple-step connections, binding-constructs specifications can be given for XML and relational data models, among others. For example, in the relational model, locations could be represented using relations, tuples, and values, and links could be represented using foreign keys. Similarly, for XML, locations could be represented using elements, element definitions, pcd data, and cdata, and links could be represented using attributes and subelement relationships. As future work, we intend to define navigation for additional data models using binding-constructs specifications.

Finally, we believe the approach described in this chapter provides a simple, generic abstraction (that consists of links and locations) that can be applied over arbitrary data models. More than that, with the binding-constructs approach, it becomes relatively straightforward to specify the links and locations for a data model, thus enabling generic and uniform access to information represented in any underlying data model (represented in the ULD). We believe that this approach can be extended beyond navigation to include querying information sources (that is, querying links and locations) and for specifying high-level mappings between data sources.

## Chapter 6

# Conclusion

Data models provide the basic data structures used to organize and store data within an information management system. Data models differ in their choice of structures, constraints, and schema requirements. The use of different data models results in structural heterogeneity, which can also include differences in schema (often due to the use of heterogeneous data models) and in how data values are represented. Structural heterogeneity must be resolved to combine information from multiple sources and to enable tool interoperability.

Most approaches for managing structural heterogeneity [1, 10, 24, 35, 40, 46, 57, 62, 64, 68, 69, 70, 75] focus on differences between schemas and assume a common data model is used for schema mapping and integration. Alternatively, a meta-data-model can be used to explicitly represent differences in data models. Many meta-data-model approaches [5, 6, 8, 31, 32, 78, 30, 63] provide languages for converting schema and corresponding instance data in one source to valid schema and instance data in another source (thus supporting one style of data-model mapping). These approaches, however, are characterized by their assumption that schema is required by all data models. Because of this assumption, these approaches cannot accurately describe more flexible data models, such as XML and RDF, cannot capture the various conformance constraints inherent in data models, and cannot express many of the desired mappings between sources that use agreed-upon conventions for representation or that must leverage multiple levels of information in a source.

The Uni-Level Description (ULD) extends existing meta-data-models by separating schema and instance constructs (through three distinct instance-of relationships) and by providing uniform access to data model, schema, and instance data. The ULD can describe

a wide variety of data models including those that allow optional, partial, and multiple levels of schema, and can be used to describe inherent data-model constraints. Finally, the ULD enables flexible transformations among information sources due to its uniform representation of data model, schema, and data.

In the rest of this chapter, we discuss our experiences with implementing the ULD for transformation and navigation, compare the ULD with other meta-data-model approaches, highlight the contributions of the ULD, and discuss possible directions for future research.

## 6.1 Implementation Issues

The ULD framework evolved from our earlier work [16, 17, 19, 20, 37], which uses a simpler meta-data-model primarily consisting of record-based structures (*struct-ct* constructs), atomic structures (*atomic-ct* constructs), and union types (*union-ct* constructs). Much of our implementation work centers on this earlier meta-data-model. Here, we briefly discuss implementations for transformation and navigation and a possible extension that we leave for future work.

For both transformation and navigation, we treat a ULD configuration as a fully materialized view [71] over an underlying information source. That is, we convert schema and data of an information source (such as an XML document) into a corresponding set of ULD facts. We assume the data-model description is also included in each configuration. For each tool that uses a particular data model, a program is written called a ULD *extractor*, which accesses the information source of interest to create a new configuration loaded with the appropriate set of ULD facts. Along with a ULD extractor, a ULD *injector* is written (for transformation) that takes a ULD configuration and creates a new information source (for example, a new XML file or relational database).

The implementations for both transformation and navigation use the Prolog programming language. (We also experimented with an early version of transformation written in Java, that stored ULD configurations in a relational database.) Using Prolog, a ULD configuration is stored as a knowledge-base, which is a text file that contains a set of facts and (possibly) Horn-clause rules. Figure 6.1 shows the basic transformation architecture.




Figure 6.1: The basic architecture of the ULD transformation engine.

(The XML configuration is assumed to be the source of the transformation and the relational configuration is assumed to be the target.) In particular, we use SWI-Prolog [81], which is an open-source Prolog implementation, to execute a set of mapping rules. The transformation engine (a Prolog program) assumes there is exactly one source and one target configuration. Both the source and target configurations serve as input to the transformation engine.

Because Prolog permits only a single formula in the head of a rule, we use a special unary predicate (that takes a list of ULD formulas) to simulate the use of multiple formulas. The transformation engine first creates a new configuration to hold the result of the transformation and updates this new configuration with facts from the given target configuration. The transformation engine then uses built-in Prolog predicates to find all answers for the head formula (used to wrap multiple target formulas). The target formulas generated from the transformation are then added to the result configuration. We note that mapping rules must be written to mimic Datalog rules, for example, to ensure that all results are obtained (using cut and fail operators), and to ensure proper stratification (by manually ordering mapping rules).

Although writing mapping rules is more difficult in Prolog (because it does not explicitly compute a fixed-point), it offers some benefits. For example, in Prolog it is straightforward to write user-defined functions and most Prolog systems provide a number of built-in libraries with functions that can be used directly in mapping rules (for example, for string manipulation). Also, the SWI-Prolog system has a built-in parser for accessing XML and XML-DTD documents, RDF and RDF-Schema documents, and it supports access to relational databases through ODBC. Thus, we can write ULD extractor and injector operations (for relational, XML, and RDF sources) directly in Prolog.

Similar to transformation, applications for data-model-aware and an early version of direct-browsing are written in Prolog [18]. Figure 6.2 shows the basic architecture for the navigation implementation. We note that navigation does not require updating an information source, that is, the architecture of Figure 6.2 only requires tools for extracting underlying information into a ULD configuration and not for injecting information back to the underlying source. For browsing, we use Amzi logic server [4] (another Prolog implementation) because it can be accessed from within Visual Basic—the implementation language of the graphical interface. Thus, similar to transformation, operator binding rules are executed using Prolog.

We note that the motivation for incremental navigation, including an initial implementation, came from our work on the Superimposed-Schematics browser [21], which uses similar operators as those defined in Chapter 5 for navigation.

A potential disadvantage of using fully materialized ULD configurations is that all information in a source is initially converted to the ULD, regardless of whether it is needed for transformation or browsing. Figure 6.3 shows a modified transformation architecture in which ULD configurations are only partially materialized with data-model information. (Note that this architecture has not been implemented.) As shown, a ULD configuration in this new architecture becomes an interface that interacts directly with the underlying source. When certain facts are needed in a transformation, the transformation engine calls functions defined in the ULD interface, which then issues queries to the underlying source, converts the result to corresponding ULD facts, and returns the associated facts to the transformation engine.




Figure 6.2: The basic architecture of the ULD navigation prototype.

The ULD interface (as shown in Figure 6.3) would contain one or more operations for each ULD formula. For example, we might have two operations for the  $ct-inst(x,y)$  formula, one to retrieve constructs  $x$  given a construct-type  $y$ , and one to retrieve a construct type  $y$  given a construct  $x$ . The ULD interface would be implemented for each tool that uses a specific data model. (Note that we can use SWI-Prolog directly to implement an interface for relational, XML, and RDF sources.) The ULD interface must also consider update operations to the underlying source for transformation.

There are a number of open issues in terms of the benefits and feasibility of using a ULD interface and partially materialized configurations. For example, one issue concerns the use of identifiers within the ULD. In the relational data model, a table construct is defined as a bag of tuple constructs. Instances of table and tuple constructs are assigned unique identifiers, which are constructed when a ULD configuration is materialized. In practice, a relational DBMS may or may not provide access to tuple identifiers. Therefore, the




Figure 6.3: An architecture for ULD transformation using non-materialized views.

ULD interface for the relational data model must assign tuple identifiers and potentially keep track of which tuple identifiers belong to underlying tuples. Finally, we note that the use of partially materialized configurations is for efficiency. There are a number of other possibilities for providing greater efficiency, for example, using existing Datalog optimization techniques or developing techniques specifically for the ULD transformation language. In general, there are many opportunities for optimizing data-model and schema transformations.

## 6.2 Related Work

Here we summarize specific meta-data-model approaches related to the ULD. In general, the ULD is the only meta-data-model approach we know of that explicitly models conformance. We believe this is due, in part, to the focus of existing approaches on solving database interoperability problems, thus assuming schema is required. In addition, in most approaches, transformation is defined at the data model or schema level (but not both). The ULD enables mappings at various levels of abstraction, for example, data-model-to-data-model, data-model-to-schema, and various combinations of these. We also

permit mappings that do not require complete, or one-to-one, transformations.

Atzeni and Torlone [5, 6, 78] define a meta-data-model (called MDM) to support mappings among data models to convert schemas in one data model into valid schemas in another data model. In their early work [5], they focus on eliminating the need to manually map among data models by using a *super data model*. For each data model  $A$ , two mappings are defined: (1) the constructs of  $A$  are mapped to the constructs of the super data model, and (2) the constructs of the super data model are mapped to the constructs of  $A$ . Thus, to convert a schema from data model  $A$  to  $B$ , these mappings are used to first convert the schema to the super data model, and then from the super data model to  $B$ . This approach avoids the need to map every data model to every other data model, that is, schemas from  $A$  can be converted to valid schemas of  $B$  without a direct data-model mapping from  $A$  to  $B$ .

MDM uses a set of basic structural primitives to define data-model constructs. The structural primitives are similar to E-R entity and relationship types and constructs are defined as compositions of these structures. For example, to describe an object-oriented data model, we would use the basic primitives to define the structure of classes and attributes. A schema in MDM is an instantiation of these data-model constructs. It is important to note that MDM does not consider source data. Thus, a schema type is a ground instance in MDM (that is, it cannot be further instantiated). Atzeni and Torlone also define approaches to automatically map data models based on an early meta-data-model, however, their approach assumes constructs and basic structures are one-to-one in a given data model (although structures can still be nested).

YAT [30, 32] differs from MDM in that it uses a meta-data-model to support data conversion (as opposed to schema conversion) among information sources. YAT is closely tied to XML, and is based on a hierarchical data model in which the data of an information source is represented as a tree, where nodes contain either labels or content and edges represent parent-child connections (attributes are not considered). YAT extends XML by offering a meta-data-model for defining *tree patterns*, which are used to create abstract DTDs (that is, DTDs with variables). These abstract DTDs describe the conventions (that is, tag names and tree structures) used to store information from other data models

in YAT. A tree pattern allows variables as nodes and cardinalities on edges. Instance data in YAT is called a *ground pattern* because it contains constants only.

YAT uses a declarative, Horn-clause language called YATL to define schema mappings for data conversion. In YATL, rules define transformations between ground trees. However, it is not possible in YATL to define transformations at higher levels of abstraction, for example, to define mappings among data models.

When compared to MDM, YAT has a limited capability for defining the structure of data model constructs or schema. YAT is focused on defining conventions for encoding heterogeneous data as hierarchies, and does not provide an explicit description of the structures commonly used in data models.

Gangopadhyay and Barsalou [8] define a meta data model called M(DM) to support schema mapping and data conversion. M(DM) introduces a set of *metatypes*, which are similar to Atzeni and Torlone's primitive structures in that metatypes denote entities (called *things*) and relationships (called *links*). A data model in M(DM) is represented as a collection of metatypes, some of which may be introduced by the data model as specializations of default M(DM) metatypes. In M(DM), a data model construct is one-to-one with a metatype as opposed to a pattern built of basic structures as in MDM or YAT. A schema instantiates the associated data model metatypes, and a database contains instances of the corresponding schema types. Thus, metatypes are considered second-order types.

The primary goal of M(DM) is to enable data intergration in federated database systems. M(DM) provides a language called M(DM)-L for defining Horn-clause views over source schemas. M(DM) can use metatypes to uniformly access source schemas in different data models for defining global views. Thus, transformation of a source to a canonical data model is not required. However, in M(DM) the structure of a metatype is not explicit: There is no mechanism in M(DM) to range over or access the metatype definitions or structures. Therefore, M(DM) cannot support data-model mappings. We note that M(DM)-L has a similar goal as YATL, but is fundamentally different because it can use structure to access schema directly, whereas YAT-L defines views as tree transformations.

The Meta-Object Facility (MOF) standard [61] is a meta-data-model architecture proposed by the Object Management Group. The MOF is similar to MDM, however, it contains a complex and large number of primitive structures (inspired by the UML meta-model) for defining data models, called *metamodels* in MOF. The main focus of the MOF is to enable information models, such as UML class diagrams, to be shared between applications. One application of the MOF architecture is to store and interchange UML class diagrams using the XML Metadata Interchange (XMI). Given a UML model, XMI prescribes a method to encode the model as an XML DTD. In addition, the MOF has an implementation in CORBA using meta objects, for example, for describing IDL interfaces. One of the primary goals of MOF is to permit uniform access to metamodels (such as UML and IDL) and models (for example, UML class diagrams).

Finally, no other meta-data-model approach we are aware of defines a constraint language for expressing inherent data-model constraints. In the ULD, constraints can define structural restrictions (such as primary and foreign key constraints) as well as requirements on conformance. Because the ULD constraint language can be represented as a stratified Datalog $\neg$  program, data-model constraints can also be executed to check that a configuration is correct (with respect to the given constraints). We believe, as discussed in Section 4.3, that constraints may also be used to check whether executing transformation rules will result in consistent configurations. In general, we believe that the ULD's treatment of conformance and constraints can be incorporated in and benefit current meta-data-model approaches, such as MOF.

### 6.3 Contributions

A key contribution of the ULD framework is its use of three, distinct *instance-of* relationships plus a conformance relationship. To the best of our knowledge, no other data-modeling or meta-data-modeling approach distinguishes type-instance relationships in this way. That is, the traditional *instance-of* or *type* relationship is not overloaded in the ULD compared with other meta-data-model approaches and data models.

The conformance relationship allows the configurer of a data model to specify the

connection between a data construct definition with the corresponding schema construct definition. And by providing constraints on the conformance relation, the ULD permits accurate description of a wide range of existing data models, including traditional database data models (with strict schema requirements) and more flexible data models such as XML and RDF.

These features of the ULD framework form the basis of the contributions of this dissertation, which we highlight below.

- *A uniform representation.* The use of distinct *instance-of* relationships allows construct types, constructs, and construct instances to be treated uniformly. Thus, data model, schema, and instance data of an information source can be directly accessed, for example, within a single query, constraint, or mapping.
- *A first-order logic interpretation.* The ULD has a direct interpretation in first-order logic, and in particular, a configuration can be represented as a set of Datalog facts and accessed using standard Datalog query languages, including stratified Datalog $\neg$ .
- *An expressive constraint language.* The ULD constraint language provides a convenient syntax for expressing constraints. The constraint language is a subset of full first-order logic, but is suitable for expressing most constraints for the data models we consider. In addition, the constraint language has a direct implementation using standard stratified Datalog $\neg$ .
- *A language for expressing common data-model constraints.* Constraint macros offer a concise syntax for expressing many of the common constraints inherent in data models. Constraint macros have a direct conversion to equivalent expressions in the ULD constraint language.
- *A classification of conformance properties.* The ULD clarifies the schema requirements offered by existing data models by allowing these requirements to be explicitly modeled. These various schema-instance arrangements are used to define a classification scheme for the structural and process-oriented conformance properties of data models.

- *A flexible transformation language.* The ULD transformation language is based on stratified Datalog $\neg$  and allows declarative, executable specifications of a wide variety of transformations. The transformation language can express mapping rules to convert information from one source to another, in the presence of different schemas and data models, thus providing a mechanism to resolve structural heterogeneity.
- *A generic language for incremental navigation.* The ULD offers the potential of applying generic operations—operations that are independent of a particular data model. We introduce generic operations for incremental navigation, which provide a high-level abstraction of links and locations for navigating and browsing an information source. With the ULD, we are able to implement the navigation operations using either ULD queries or high-level binding specifications so that the operations can be easily applied generically to information sources regardless of the data model used for representation.

Finally, in this dissertation, we used the ULD to: (1) express the constructs and constraints of a wide variety of data models, with varying structures and conformance properties, (2) define a number of common data-model transformations as well as a schema transformation that resolves structural heterogeneity, (3) sketch approaches to validate transformation rules and fill in partial rules, and (4) specify data-model aware and direct navigation bindings for the RDF data model.

## 6.4 Future Work

We describe areas of future work beyond the optimization issues discussed earlier. The following open issues are focused on transformation, and specifically on making it easier to define mappings to resolve structural heterogeneity and to further understand and exploit the details of data models and their constraints.

One reason writing transformations is difficult is that data models and schemas have various implicit constraints. Keeping track of these constraints and ensuring that a transformation does not violate them is generally difficult. Although the ULD provides an expressive constraint language (so that constraints can be explicitly captured), not all of

the constraints that can be expressed can be exploited. In particular, checking that a transformation is satisfiable in the presence of constraints (that are expressed using stratified Datalog $\neg$ ) is in general an undecidable problem. The constraint macros described in Chapter 3 offer one approach to limit the set of possible constraints (possibly to a decidable set). Of interest is to define a broad enough set of constraint macros to capture important data-model restrictions (such as conformance constraints), while maintaining the decidability of checking for well-formed transformations. A similar approach is used in description logics [7], where the goal is to define a decidable fragment of first order logic for constraining semantic data models. As in description logics, our challenge is to define classes of constraints that can usefully restrict data models while at the same time be practically exploited for transformation.

The ULD transformation language specifies the low-level details of a conversion. Such an approach is attractive because it enables extremely flexible transformations (for example, data-model mappings, schema mappings, and arbitrary combinations), however, specifying mapping rules is not trivial. We would like to explore the use of higher-level transformation languages that shield the low-level details of transformation. Chapter 4 sketched an informal approach for elaborating constraints, which assumes an initial correspondence between data-model constructs. Similar, correspondence-based approaches have been defined for schema mappings [64, 68, 69]. We would like to look further into applying these approaches for data-model mappings. The ULD might also provide benefits for existing schema-based approaches. For instance, if a mapping among schemas can be expressed in the ULD transformation language given initial schema correspondences, the data-model mapping would result “for free”—that is, with the ULD, we bypass the common data model, eliminating the additional conversions in and out of the common data model from source and target data models.

We also want to explore the use of transformation patterns for expressing mappings. A transformation pattern captures common mappings that can be combined to form complex structural conversions. Although not described here, we have experimented with transformation patterns [19] as parameterized specifications that take simple correspondences as input to generate complex transformations. In this approach, patterns are analogous

to higher-order functions like map and fold in functional programming languages.

Another area of interest concerns modeling standard conventions with the ULD. Most data models have well-known modeling conventions for representing data structures that cannot be expressed directly with the data model's available constructs. For example, in the relational model, there are no explicit constructs for representing hierarchy; however, there are many standard conventions (or strategies) for storing hierarchical data using relations (some of these were used in the XML to relational transformations of Chapter 4). We believe the ULD can be used to describe such modeling conventions (perhaps as transformation patterns). By capturing modeling conventions and relating them to standard modeling concepts in conceptual modeling theory, it may be possible to automate data-model transformations. In addition, such a repository of conventions could serve as a guide for database design.

# Bibliography

- [1] ABITEBOUL, S., CLUET, S., AND MILO, T. Correspondence and translation for heterogeneous data. In *Proceedings of the 6th International Conference on Database Theory (ICDT)* (1997), vol. 1186 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 351–363.
- [2] ABITEBOUL, S., HULL, R., AND VIANU, V. *Foundations of Databases*. Addison-Wesley Publishing Company, 1995.
- [3] ALBERT, J. Theoretical foundations of schema restructuring in heterogeneous multidatabase systems. In *Proceedings of the ACM CIKM International Conference on Information and Knowledge Management* (2000), ACM Press, pp. 461–470.
- [4] AMZI! INC. *Amzi! Prolog + Logic Server*. Available: <http://www.amzi.com/>, [Viewed: December 11, 2003].
- [5] ATZENI, P., AND TORLONE, R. Schema translation between heterogeneous data models in a lattice framework. In *Proceedings of the 6th IFIP TC-2 Working Conference on Data Semantics (DS-6)* (1995), Chapman and Hall, pp. 345–364.
- [6] ATZENI, P., AND TORLONE, R. Management of multiple models in an extensible database design tool. In *Proceedings of the 5th International Conference on Extending Database Technology (EDBT)* (1996), vol. 1057 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 79–95.
- [7] BAADER, F., CALVANESE, D., MCGUINNESS, D., NARDI, D., AND PATEL-SCHNEIDER, P., Eds. *The Description Logic Handbook*. Cambridge University Press, 2003.
- [8] BARSALOU, T., AND GANGOPADHYAY, D. M(DM): An open framework for interoperation of multimodel multidatabase systems. In *Proceedings of the 8th International Conference on Data Engineering (ICDE)* (1992), IEEE Computer Society, pp. 218–227.
- [9] BARTA, R. AsTMa\* language family. Tech. rep., Bond University, July 2003. Available: <http://astma.it.bond.edu.au/astma-family.dbk>, [Viewed: December 11, 2003].

- [10] BATINI, C., LENZERINI, M., AND NAVATHE, S. B. A comparative analysis of methodologies for database schema integration. *ACM Computing Surveys* 18, 4 (1986), 323–364.
- [11] BERNERS-LEE, T., HENDLER, J., AND LASSILA, O. The semantic web. *Scientific American* 284, 5 (2001), 34–43.
- [12] BERNSTEIN, P. A., HALEVY, A. Y., AND POTTINGER, R. A vision of management of complex models. *SIGMOD Record* 29, 4 (Dec. 2000), 55–63.
- [13] BIEZUNSKI, M., BRYAN, M., AND NEWCOMB, S. R., Eds. *Topic Maps*. ISO/IEC 13250: Information Technology – Document Description and Markup Languages, December 1999. Available: <http://www.y12.doe.gov/sgml/sc34/document/0129.pdf>, [Viewed: December 11, 2003].
- [14] BOHANNON, P., FREIRE, J., ROY, P., AND SIMÉON, J. From XML Schema to relations: A cost-based approach to XML storage. In *Proceedings of the 18th International Conference on Data Engineering (ICDE)* (2002), IEEE Computer Society, pp. 64–73.
- [15] BOOCH, G., RUMBAUGH, J., AND JACOBSON, I. *The Unified Modeling Language Users Guide*. Addison-Wesley Publishing, 1998.
- [16] BOWERS, S., AND DELCAMBRE, L. Representing and transforming model-based information. In *Electronic Proceedings of the ECDL International Workshop on the Semantic Web* (2000), pp. 1–14. Available: <http://www.ics.forth.gr/isl/SemWeb/proceedings/session1-1/paper.pdf>, [Viewed: December 11, 2003].
- [17] BOWERS, S., AND DELCAMBRE, L. A generic representation for model-based information. *Electronic Transactions on Artificial Intelligence (ETAI) [Online]* 5, 4 (2001), 1–34. Available: <http://www.ida.liu.se/ext/epa/ej/etai/2001/D/index.html>, [Viewed: December 11, 2003], printed version published by *Royal Swedish Academy of Sciences*.
- [18] BOWERS, S., AND DELCAMBRE, L. JustBrowsing: A generic API for exploring information. In *Demo Session at the 21st International Conference on Conceptual Modeling (ER)* (2002).
- [19] BOWERS, S., AND DELCAMBRE, L. On modeling conformance for flexible transformation over data models. In *Knowledge Transformation for the Semantic Web* (2003), vol. 95 of *Frontiers in Artificial Intelligence and Applications*, ISO Press, pp. 34–48.

- [20] BOWERS, S., AND DELCAMBRE, L. The uni-level description: A uniform framework for representing information in multiple data models. In *Proceedings of the 22nd International Conference on Conceptual Model (ER)* (2003), vol. 2813 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 45–58.
- [21] BOWERS, S., DELCAMBRE, L., AND MAIER, D. Superimposed schematics: Introducing E-R structure for *in-situ* information selections. In *Proceedings of the 21st International Conference on Conceptual Model (ER)* (2002), vol. 2503 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 90–104.
- [22] BRAY, T., PAOLI, J., SPERBERG-MCQUEEN, C., AND MALER, E., Eds. *Extensible Markup Language (XML) 1.0 (Second Edition)*. W3C Recommendation. World Wide Web Consortium (W3C), October 2000. Available: <http://www.w3.org/TR/2000/REC-xml-20001006>, [Viewed: December 11, 2003].
- [23] BRICKLEY, D., AND GUHA, R., Eds. *RDF Vocabulary Description Language 1.0: RDF Schema*. W3C Working Draft. World Wide Web Consortium (W3C), February 2003. Available: <http://www.w3.org/TR/2003/WD-rdf-schema-20030123>, [Viewed: December 11, 2003].
- [24] CALI, A., CALVANESE, D., GIACOMO, G. D., AND LENZERINI, M. On the expressive power of data integration systems. In *Proceedings of the 21st International Conference on Conceptual Modeling (ER)* (2002), vol. 2503 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 338–350.
- [25] CAREY, M. J., HAAS, L. M., MAGANTY, V., AND WILLIAMS, J. H. PESTO: An integrated query/browser for object databases. In *Proceedings of 22nd International Conference on Very Large Data Bases (VLDB)* (1996), Morgan Kaufmann, pp. 203–214.
- [26] CATTELL, R., Ed. *The Object Database Standard: ODMG 2.0*. Morgan Kaufmann Publishers, 1997.
- [27] CHEN, P. P. The entity-relationship model—toward a unified view of data. *ACM Transactions on Database Systems (TODS)* 1, 1 (1976), 9–36.
- [28] CHEN, P. P. A preliminary framework for Entity-Relationship models. In *Proceedings of the National Computer Conference* (1977), AFIPS Press, pp. 77–84.
- [29] CHRISTOPHIDES, V., ABITEBOUL, S., CLUET, S., AND SCHOLL, M. From structured documents to novel query facilities. In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data* (1994), ACM Press, pp. 313–324.

- [30] CHRISTOPHIDES, V., CLUET, S., AND SIMÉON, J. On wrapping query languages and efficient xml integration. In *Proceedings of the ACM SIGMOD International Conference on Management of Data* (2000), ACM Press, pp. 141–152.
- [31] CLAYPOOL, K. T., AND RUDENSTEINER, E. A. Sangam: A framework for modeling heterogeneous database transformations. In *Proceedings of the 5th International Conference on Enterprise Information Systems* (2003), vol. 1, pp. 219–224.
- [32] CLUET, S., DELOBEL, C., SIMÉON, J., AND SMAGA, K. Your mediators need data conversion! In *Proceedings of the 1998 ACM SIGMOD International Conference on Management of Data* (1998), ACM Press, pp. 177–188.
- [33] CODD, E. A relational model of data for large shared data banks. *Communications of the ACM (CACM)* 13, 6 (1970), 377–387.
- [34] COHEN, W. W. Some practical observations on integration of web information. In *Informal Proceedings of the ACM SIGMOD Workshop on the Web and Databases (WebDB)* (1999), pp. 55–60.
- [35] DAVIDSON, S. B., AND KOSKY, A. WOL: A language for database transformations and constraints. In *Proceedings of the 13th International Conference on Data Engineering (ICDE)* (1997), IEEE Computer Society, pp. 55–65.
- [36] DELCAMBRE, L., AND MAIER, D. Models for superimposed information. In *Advances in Conceptual Modeling (ER)* (1999), vol. 1727 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 264–280.
- [37] DELCAMBRE, L., MAIER, D., BOWERS, S., WEAVER, M., DENG, L., GORMAN, P., ASH, J., LAVELLE, M., AND LYMAN, J. Bundles in captivity: An application of superimposed information. In *Proceedings of the 17th International Conference on Data Engineering (ICDE)* (2001), IEEE Computer Society, pp. 111–120.
- [38] ELMASRI, R., AND NAVATHE, S. B. *Fundamentals of Database Systems*. Addison-Wesley Publishing Company, 1994.
- [39] GENESERETH, M. R., AND NILSSON, N. J. *Logic Foundations of Artificial Intelligence*. Morgan Kaufmann Publishers, Inc., 1987.
- [40] HAMMER, J., AND MCLEOD, D. On the resolution of representational diversity in multidatabase systems. In *Management of Heterogeneous and Autonomous Database Systems*. Morgan Kaufmann, 1998, pp. 91–118.

- [41] HAYES, P., AND MCBRIDE, B., Eds. *RDF Semantics*. W3C Working Draft. World Wide Web Consortium (W3C), January 2003. Available: <http://www.w3.org/TR/2003/WD-rdf-mt-20030123>, [Viewed: December 11, 2003].
- [42] HENDERSEN-SELLERS, B. Advanced OO modelling: Metamodels and notations for the new millennium. In *Proceedings of the 21st International Conference on Conceptual Modeling (ER)* (2002), vol. 2503 of *Lecture Notes in Computer Science*, Springer-Verlag, p. 11.
- [43] JARKE, M., EHERER, S., GALLERSDÓRFER, R., JEUSFELD, M. A., AND STAUDT, M. ConceptBase – A deductive object base manager. Tech. rep., Aachener Informatik-Berichte 93-14, 1995.
- [44] KICZALES, G., RIVIERES, J. D., AND BOBROW, D. *The Art of the Metaobject Protocol*. MIT Press, 1991.
- [45] KIFER, M., KIM, W., AND SAGIV, Y. Querying object-oriented databases. In *Proceedings of the ACM SIGMOD International Conference on Management of Data* (1992), ACM Press, pp. 393–402.
- [46] KRISHNAMURTHY, R., LITWIN, W., AND KENT, W. Language features for interoperability of databases with schematic discrepancies. In *Proceedings of the 1991 ACM SIGMOD International Conference on Management of Data* (1991), ACM Press, pp. 40–49.
- [47] LASSILA, O., AND SWICK, R. R., Eds. *Resource Description Framework (RDF) Model and Syntax Specification*. W3C Recommendation. World Wide Web Consortium (W3C), February 1999. Available: <http://www.w3.org/TR/1999/REC-rdf-syntax-19990222>, [Viewed: December 11, 2003].
- [48] LAUSEN, G., AND VOSSEN, G. *Models and Languages of Object-Oriented Databases*. Addison-Wesley, 1998.
- [49] LÉCLUSE, C., RICHARD, P., AND VÉLEZ, F. O<sub>2</sub>, an object-oriented data model. In *Proceedings of the 1988 ACM SIGMOD International Conference on Management of Data* (1988), ACM Press, pp. 424–433.
- [50] LEVY, A. Y., AND ROUSSET, M.-C. Verification of knowledge bases based on containment checking. *Artificial Intelligence* 101, 1-2 (1998), 227–250.
- [51] LEVY, A. Y., AND SAGIV, Y. Constraints and redundancy in Datalog. In *Proceedings of the Eleventh ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems* (1992), ACM Press, pp. 67–80.

- [52] LUDÄSCHER, B., GUPTA, A., AND MARTONE, M. E. Model-based mediation with domain maps. In *Proceedings of the 17th International Conference on Data Engineering (ICDE)* (2001), IEEE Computer Society, pp. 81–90.
- [53] LUGER, G. F. *Artificial Intelligence: Structures and Strategies for Complex Problem Solving (4th Edition)*. Addison-Wesley Publishing, 2002.
- [54] MARK, L., AND ROUSSOPOULOS, N. Integration of data, schema and meta-schema in the context of self-documenting data models. In *Proceedings of the 3rd International Conference on Entity-Relationship Approach (ER)* (1983), North-Holland, pp. 585–602.
- [55] MARSHALL, C. C., SHIPMAN III, F. M., AND COOMBS, J. H. VIKI: Spatial hypertext supporting emergent structure. In *European Conference on Hypertext Technology (ECHT)* (1994), ACM Press, pp. 13–23.
- [56] MCGUINNESS, D. L., AND VAN HARMELEN, F., Eds. *OWL Web Ontology Language Overview*. W3C Candidate Recommendation. World Wide Web Consortium (W3C), August 2003. Available: <http://www.w3.org/TR/2003/CR-owl-features-20030818>, [Viewed: December 11, 2003].
- [57] MILLER, R. J. Using schematically heterogeneous structures. In *Proceedings of the ACM SIGMOD International Conference on the Management of Data* (1998), ACM Press, pp. 189–200.
- [58] MYCROFT, A., AND O’KEEFE, R. A. A polymorphic type system for Prolog. *Artificial Intelligence* 23, 3 (1984), 295–307.
- [59] MYLOPOULOS, J., BORGIDA, A., JARKE, M., AND KOUBARAKIS, M. Telos: Representing knowledge about information systems. *ACM Transactions on Information Systems* 8, 4 (1990), 325–362.
- [60] NOY, N., FERGERTSON, R., AND MUSEN, M. The knowledge model of Protege-2000: Combining interoperability and flexibility. In *Proceedings of the 2nd International Conference on Knowledge Engineering and Knowledge Management (EKAW)* (2000), vol. 1937 of *Lecture Notes in Computer Science*, Springer-Verlag.
- [61] OMG. *Meta Object Facility (MOF) Specification*, September 1997. OMG Document ad/97-08-14, Available: <http://www.dstc.edu.au/Research/Projects/MOF/Publications.html>, [Viewed: December 11, 2003].

- [62] PAPAKONSTANTINOY, Y., ABITEBOUL, S., AND GARCIA-MOLINA, H. Object fusion in mediator systems. In *Proceedings of 22nd International Conference on Very Large Data Bases (VLDB)* (1996), Morgan Kaufmann, pp. 413–424.
- [63] PAPAZOGLU, M. P., AND RUSSELL, N. A semantic meta-modeling approach to schema transformation. In *Proceedings of the International Conference on Information and Knowledge Management (CIKM)* (1995), ACM Press, pp. 113–121.
- [64] PARENT, C., AND SPACCAPIETRA, S. Issues and approaches of database integration. *Communications of the ACM* 41, 5 (May 1998), 166–178.
- [65] PEPPER, S., Ed. *Requirements for a Topic Map Constraint Language (Draft)*. ISO/IEC JTC 1/SC34: Information Technology – Document Description and Markup Languages, April 2003. Available: <http://www.isotopicmaps.org/tmcl/requirements.html>, [Viewed: December 11, 2003].
- [66] PEPPER, S., AND MOORE, G., Eds. *XML Topic Maps (XTM) 1.0*. TopicMaps.Org Authoring Group, Specification, August 2001. Available: <http://www.topicmaps.org/xtm/1.0/xtm1-20010806.html>, [Viewed: December 11, 2003].
- [67] POPA, L., AND TANNEN, V. An equational chase for path-conjunctive queries, constraints, and views. In *Proceedings of the 7th International Conference on Database Theory (ICDT)* (1999), vol. 1540 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 39–57.
- [68] POPA, L., VELEGRAKIS, Y., MILLER, R. J., HERNÁNDEZ, M., AND FAGIN, R. Translating Web data. In *Proceedings of the 28th International Conference on Very Large Data Bases (VLDB)* (2002), Morgan Kaufmann, pp. 598–609.
- [69] POTTINGER, R., AND BERNSTEIN, P. A. Merging models based on given correspondences. In *Proceedings of the 29th International Conference on Very Large Data Bases (VLDB)* (2003), Morgan Kaufmann, pp. 826–837.
- [70] RAHM, E., AND BERNSTEIN, P. A. A survey of approaches for automatic schema matching. *The VLDB Journal* 10, 4 (2001), 334–350.
- [71] RAMAKRISHNAN, R., AND GEHRKE, J. *Database Management Systems, 3rd Edition*. McGraw Hill, 2003.

- [72] RATH, H. H. The topic maps handbook. Tech. rep., empolis GmbH, 2003. Available: [http://www.empolis.com/download/docs/whitepapers/empolistopicmapswhitepaper\\_eng.pdf](http://www.empolis.com/download/docs/whitepapers/empolistopicmapswhitepaper_eng.pdf), [Viewed: December 11, 2003].
- [73] REITER, R. Towards a logical reconstruction of relational database theory. In *On Conceptual Modeling*, M. L. Brodie, J. Mylopoulos, and J. W. Schmidt, Eds. Springer-Verlag, 1984.
- [74] SHANMUGASUNDARAM, J., TUFTE, K., HE, G., ZHANG, C., DEWITT, D., AND NAUGHTON, J. Relational databases for querying XML documents: Limitations and opportunities. In *Proceedings of the 25th International Conference on Very Large Data Bases (VLDB)* (1999), Morgan Kaufmann, pp. 302–314.
- [75] SHETH, A. P., AND LARSON, J. A. Federated database systems for managing distributed, heterogeneous, and autonomous databases. *ACM Computing Surveys* 22, 3 (1990), 183–236.
- [76] SHIPMAN III, F. M., WEI HSIEH, H., MALOOR, P., AND MOORE, J. M. The visual knowledge builder: A second generation spatial hypertext. In *The 12th ACM Conference on Hypertext and Hypermedia (Hypertext 2001)* (2001), ACM Press, pp. 113–122.
- [77] TIAN, F., DEWITT, D. J., CHEN, J., AND ZHANG, C. The design and performance evaluation of alternative XML storage strategies. *SIGMOD Record* 31, 1 (2002), 5–10.
- [78] TORLONE, R., AND ATZENI, P. A unified framework for data translation over the web. In *Proceedings of the 2nd International Conference on Web Information Systems Engineering (WISE)* (2001), IEEE Computer Society, pp. 350–358.
- [79] TSICHRITZIS, D. C., AND LOCHOVSKY, F. H. *Data Models*. Prentice-Hall, 1982.
- [80] ULLMAN, J. D. *Principles of Database and Knowledge-Base Systems, Volume II*. Computer Science Press, 1989.
- [81] UNIVERSITY OF AMSTERDAM. *SWI-Prolog*. Available: <http://www.swi-prolog.org/>, [Viewed: December 11, 2003].
- [82] WOODS, W. A. What’s in a link: Foundations for semantic networks. In *Representation and Understanding: Studies in Cognitive Science* (1975), Academic Press, pp. 35–82.

## Biographical Note

Shawn Bowers was born in Eugene, Oregon on February 18, 1973. He graduated from South-Eugene High School in 1991, and began studying Computer Science at Portland State University (PSU) in Portland, Oregon in 1993. In 1996, he transferred to the University of Oregon in Eugene, Oregon. He completed his undergraduate thesis under the direction of Professor Arthur Farley and graduated with his Bachelor of Science degree in Computer and Information Science in 1998. For six months during the Spring and Summer of 1998, Shawn worked on the Assert project studying quality-of-service validation under the direction of Professor Stephen Fickas. Shawn then returned to Portland to join the Master of Science program in Computer Science and Engineering at the Oregon Graduate Institute (OGI).<sup>1</sup>

At OGI, Shawn began working on software-component adaptation in the Heterodyne project under the direction of Professor Lois Delcambre and Professor David Maier. And with the encouragement of Lois, Shawn joined the Ph.D. program at OGI in the Fall of 1999 to study superimposed information and data models. During his Doctoral studies at OGI, Shawn interned for four months at WireX, Inc., in Portland, Oregon working with Dr. Crispin Cowan. Shawn was also an instructor several times for the Introduction to Database Systems course at OGI and PSU, and designed and taught the Artificial Intelligence course at the Oregon Institute of Technology in Beaverton, Oregon.

Currently, Shawn is a Postdoctoral Scholar at the University of California, San Diego, working in the Knowledge-Based Information Systems Lab at the San Diego Supercomputer Center.

---

<sup>1</sup>The Oregon Graduate Institute merged with the Oregon Health and Sciences University (OHSU) in 2001 and is now the OGI School of Science and Engineering at OHSU.