

# On Modeling Conformance for Flexible Transformation over Data Models\*

Shawn Bowers and Lois Delcambre

OGI School of Science & Engineering, OHSU  
Beaverton, Oregon 97006, USA  
{shawn, lmd}@cse.ogi.edu

**Abstract.** Many useful and commonly used data models exist for structured information. Such models often use only a limited number of *basic structures* to represent information such as sets, bags, lists, attribute-value pairs, and scalars. They differ, however, in whether or not they *permit* the specification of schema (or DTD or template), whether they permit *more than one* schema, and whether they *require* schema. Some even allow multiple levels of schema-instance conformance (e.g., RDF and Topic Maps) where the type of an instance is an instance, and so on. Support for transforming between such varying representation schemes remains a significant challenge. We extend our earlier work on generic representation and transformation of model-based information by introducing a richer metamodel and abstract framework for representation. We also introduce several steps toward our vision of high-level transformations where complex mapping rules are defined using transformation patterns and by exploiting inherent constraints.

## 1 Introduction

Taking information in one representation scheme (such as XML) and extracting some or all of it for use in another scheme (such as a Topic Map, a Relational database, or as RDF triples) is a surprisingly difficult task. In fact, few tools exist to help perform such transformations, which means that (potentially) complex, special purpose programs must be written for even very simple mappings. One reason such conversions are difficult is that representation schemes differ in the basic structural constructs and schema constraints they provide for organizing information. As a consequence, straightforward, one-to-one mappings between schemes rarely exist. In addition, the absence of high-level languages for expressing transformations between schemes places the burden of conversion on these special-purpose programs, making it difficult to define, maintain, and reason about transformations.

We observe, however, that structured information is typically based on a small set of structural constructs, composed in various ways. When described explicitly, these structural constructs can be exploited directly for transformation. For example, the constructs offered by the Relational model can be viewed as a set of tables, each consisting of a bag of rows, where each row has a set of named atomic or scalar attribute values. Similarly, an XML document can be viewed as a set of elements, possibly with attributes, where each element has a list of sub-elements or scalar values (PCDATA). We propose that a wide range of useful transformations can be

performed through mappings between representation scheme constructs, e.g., elements to tuples in tables, and so on.

To this end, we present a uniform framework that allows the description of arbitrary representation schemes along with a transformation language that can easily convert information within and across schemes. The goals of our research are to:

1. capture the representation scheme or data model explicitly—by instantiating and composing the basic structures (or construct types) of a metamodel,
2. support representation schemes with varying types of conformance, *i.e.*, we model schema-instance relationships for representation schemes where schema (or type information) can be required, optional, or can occur at multiple levels, and
3. express transformation rules declaratively, *e.g.*, by allowing users to indicate simple correspondences between data model constructs and converting the associated information automatically.

This paper is organized as follows. Section 2 describes various data models and representation schemes, with emphasis on how each differs in their use of conformance relationships. Section 3 presents the uniform framework for representing structured information, along with the four relationships of interest. Section 4 introduces transformation rules that exploit this representation and also discusses the possibility of higher-level transformation rules, making rule specification easier. Section 5 discusses related work and we conclude in Section 6 with a brief discussion of work in progress.

## 2 Data Models and Conformance

Database systems are *schema-first* in that a schema must be defined before any data can be placed in the database. The role of the schema is to introduce and name application-specific structures (such as the “Employee” table with “SSN” and “Name” as attributes for Employee) that will be used to hold application data. A schema imposes uniform structure and constraints on application data. Various tools such as a query processor, in turn, can exploit this uniform structure.

A number of data models exists that are not schema-first, including XML and other semi-structured models, RDF, the Topic Map Model, and various hypertext data models [1, 2]. Both RDF and XML are models where the schema, *i.e.*, the RDF schema or DTD, respectively, is optional. More than that, even with a schema, the Topic Map model, RDF, and XML (with an “open” DTD) permit additional, schema-free structures to be freely mixed with structures that conform to a schema.

We model *conformance relationships* explicitly, as appropriate, for each data model of interest. Table 1 provides

\* This work supported in part by the National Science Foundation through grants EIA 9983518 and IIS 9817492.

example data models with some of their associated structures that participate in conformance. For example, in object-oriented models, objects conform to classes, whereas in XML, elements conform to element types, and so on.

**Table 1.** Example data model structures involved in conformance

<i>Data Model</i>	<i>Selected structures, related by the conformance relationship</i>
Object-Oriented	Objects conform to Classes.
Relational	A Table of Tuples conforms to a Relation Scheme.
XML w/ DTD	Elements conform to Element Types and Attributes to Attribute Types.
RDF w/ RDFS	Objects (resources) conform to Classes and Properties to Property Types.
XTM	Topics conform to Topic Types (Topics), Associations conform to Association Types (Topics), and Occurrences conform to Occurrence Types (Topics).
E-R	Entities conform to Entity Types, Relationships conform to Relationship Types, and Values conform to Value Types.

Table 2 describes some of the properties of conformance that are present in the data models of Table 1. First, we consider how the conformance relationship is established with three possibilities: implicitly, upon request, and explicitly as shown in the first section of Table 2. In the next section of Table 2, we see that the schema-first models also require conformance for their data values whereas the other models do not. And the models with optional conformance are also open: they permit conforming and non-conforming data to be freely mixed, as shown in the third section of Table 2. Some models allow a construct to conform to several structures, namely all but the relational model (fourth section of Table 2). Finally, we see that all of the non-schema-first models except XML permit multiple levels of conformance.

To represent a wide range of data models explicitly, we must support variations in the conformance relationship as shown by Table 2. We discuss the details of our representation in the next section, discuss transformation in Section 4, and further consider related work in Section 5.

### 3 The Uniform Framework

The main component of the uniform framework is a three-level, metamodel architecture that incorporates conformance relationships. In this section, we describe our metamodel architecture, introduce a concrete metamodel (*i.e.*, a specific set of basic structures), and present a logic-based description language that enables uniform data model, schema, and instance description.

#### 3.1 The Generic Metamodel Architecture

Figure 1 exemplifies a typical metamodel architecture. As shown, the architecture has four levels: the metamodel, data model, schema, and data levels, where each item in a level is considered an instance of an item in the level above it.

**Table 2.** Examples properties of the conformance relationship

<i>Conformance Properties</i>	<i>Data Model</i>
<b>How is the conformance relationship established?</b>	
when data is created (implicitly)	Object-Oriented, Relational, and E-R
when computed (upon request)	XML w/ DTD (documents must be validated)
when conformance is declared (explicitly)	RDF w/ RDFS (via <code>rdf:type</code> property) and XTM (via class-instance association)
<b>Is conformance required?</b>	
required	Object-Oriented, Relational, and E-R
optional	RDF (RDF Schema is not required), XTM (type information not required), and XML (DTD is not required)
<b>Is conformance open?</b>	
no	Object-Oriented, Relational, and E-R
yes	RDF w/ RDFS, XTM, and XML
<b>Is multiple conformance allowed?</b>	
no	Relational
yes	Object-Oriented (inheritance), RDF w/ RDFS (Property/Class with multiple inheritance), XTM (Topics can be instances of multiple Topic Types), and E-R
<b>How many levels of conformance are permitted?</b>	
one	Relational, E-R, and Object-Oriented (when not in the next entry)
two (sometimes)	Object-Oriented (some models support classes as objects)
zero to any number	XML (zero or one), XTM, and RDF

As shown, typical metamodel architectures are characterized by their assumption that data models are schema-first, *i.e.*, schema items must be created prior to data items.

In contrast, Figure 2 shows the three-level metamodel architecture we employ. The top level (denoted as the metamodel) consists of *construct types*, which represent the basic structures used within data models for defining schema and data. Examples include types such as collection, name-value pair, atomic, and so on. The middle layer defines data model and schema *constructs* along with their composition. For example, both a `class` and `object` construct along with an explicit conformance definition between them (represented by the relationship type labeled `conf`) are defined in the middle layer. In this way, the metamodel is used to define all the constructs of the data model, not just those for creating schema.

Finally, the bottom level represents *instances* and (data-level) *conformance relationships*. For example, an object (with the value “steve”) would be connected with a particular class (with name “person”) using the `d-inst` relationship.

The architecture distinguishes three kinds of instance-of relationships, in addition to the conformance relationship, as shown in Figure 2. Constructs introduced in the middle layer are necessarily an instance (`ct-inst`, read as “construct-type instance”) of a metamodel construct type. (Note that the metamodel currently has a range of basic structures as construct types and the set of construct types can be easily extended if needed.) Similarly, any value introduced in the bottom layer is necessarily an instance (`c-inst`, read

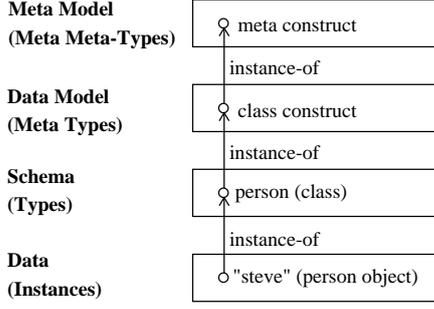


Fig. 1. A typical four-level metamodel architecture.

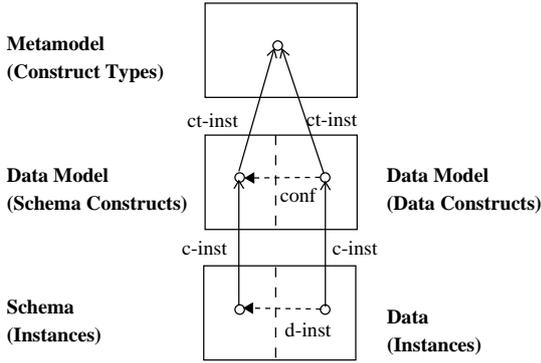


Fig. 2. The three-level metamodel architecture.

as “construct instance”) of the constructs introduced in the middle layer. The flexibility of the framework is due to the conformance relationships. Conformance relationships are expressed within the middle layer (e.g., to indicate that an XML element can optionally conform to an XML element type) and the corresponding data-level instance-of relationships (d-inst, read as “data instance”) expressed within the bottom layer (e.g., to indicate that a particular element is an instance of another element type).

To make our framework concrete, we introduce the notion of a *configuration*, which consists of a particular choice of metamodel (i.e., set of construct types), a particular data model (i.e., constructs) defined using the construct types, and a collection of instances defined using the constructs. A configuration serves as input to a transformation. Thus, a transformation takes one (or possibly more) configurations and generates a new configuration.

A configuration  $\mathcal{C}$  is a tuple  $(M, D, I, T)$  whose components are defined as follows.

- $M$  is a set of named construct types in the top layer. (*The Metamodel*)
- $D$  is a tuple  $(C, R)$  where  $C$  is a set of named constructs defined using the types of  $M$ , and  $R$  is a set of ordered pairs  $(c_1, c_2)$  for  $c_1, c_2 \in C$  denoting a conformance definition and is read as “instances of  $c_1$  can conform to instances of  $c_2$ .” (*The Middle Layer*)
- $I$  is a set of named instances. The name of an instance acts as its unique identifier. (*The Bottom Layer*)
- $T$  is a tuple  $(T_{ct}, T_c, T_d)$  in which  $T_{ct} \cup T_c \cup T_d$  is the set of type-instance relationships of the configuration where  $T_{ct}$ ,  $T_c$ , and  $T_d$  are disjoint sets of ordered pairs  $(i_1, i_2)$

such that for  $t \in T_{ct}$  we require  $i_1 \in C$  and  $i_2 \in M$  (i.e., *ct-inst*), for  $t \in T_c$  we require  $i_1 \in I$  and  $i_2 \in C$  (i.e., *c-inst*), and for  $t \in T_d$  we require  $i_1, i_2 \in I$  (i.e., *d-inst*) such that  $(i_1, c_1), (i_2, c_2) \in T_C$  and  $(c_1, c_2) \in R$ . (*Cross-Layer Connections and Data Instances*)

We informally define construct type, construct, and instance as follows.<sup>1</sup>

- A construct type  $ct$  is a tuple  $(adt, P)$  such that  $adt$  represents an abstract (or parameterized) data type and  $P$  is a set of restriction properties.
- A construct  $c$  is a tuple  $(ct, \rho)$  such that  $ct \in M$  is a construct type for which  $c$  is an instance (and hence  $(c, ct) \in T_{ct}$ ) and  $\rho$  is a set of restrictions that obey  $P$  and serve to restrict the instances of the construct and define its compositional aspects.
- An instance  $i$  is a tuple  $(c, val)$  such that  $c \in C$  and  $val$  is a valid instance of its associated construct type’s  $adt$  and obeys the restrictions  $\rho$  of  $c$ .

Intuitively, a construct type specifies a basic data structure. Thus, the value of a construct instance is an instance of its construct type’s basic data structure. A description of composition constraints on the basic structures (i.e., constructs) of a model is also given. For example, a particular construct might represent a collection of name-value pairs.

### 3.2 Example Description Language and Basic Structures

Here we introduce an example metamodel, i.e., a concrete set of metamodel construct types, and demonstrate their use for describing data models. We also present a logic-based language for representing constructs and instances.

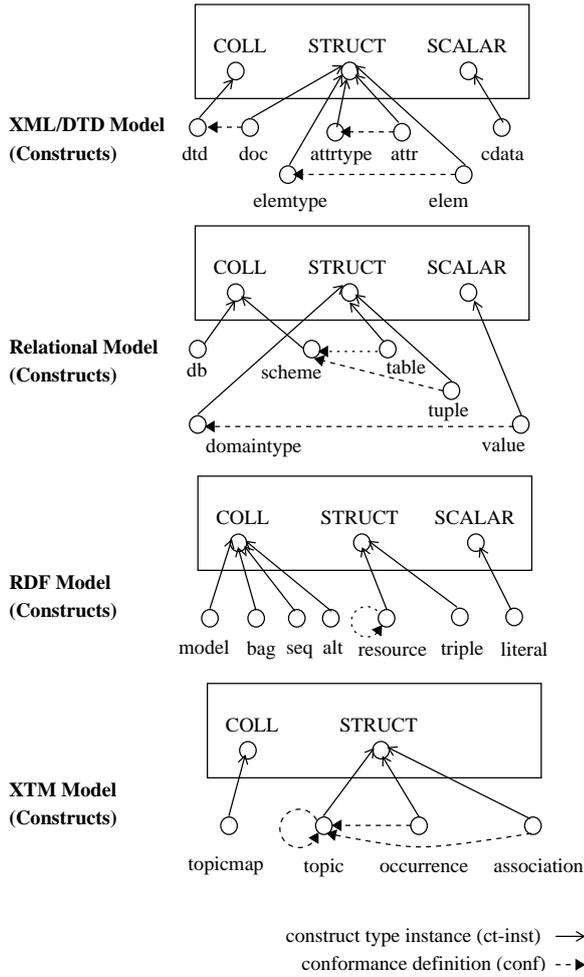
The construct types in this metamodel are collection, (specifically,  $set_{ct}$ ,  $list_{ct}$ ,  $bag_{ct}$ ),  $struct_{ct}$ , and  $scalar_{ct}$ , representing collections, name-value pairs, and atomic data (like strings and integers), respectively. We should note that this choice of construct types is only one possible set of structures that are sufficient for the data models we describe in this paper. Figure 3 shows the XML, Relational, RDF, and Topic Map data model constructs expressed using our construct types (note that we use  $COLL$  to represent the collection construct types). Conformance is shown explicitly, e.g., XML elements can conform to element types, Relational tables to schemes, RDF resources to resources, and Topic Map associations to topics.

We use the following predicates for representing the instance and conformance relationships of configurations.

```
ct-inst(c, ct).
conf(c1, c2).
c-inst(d, c).
d-inst(d1, d2).
```

The *ct-inst* predicate relates a construct (shown as  $c$ ) to a construct type (shown as  $ct$ ), which must be either  $set_{ct}$ ,  $list_{ct}$ ,  $bag_{ct}$ ,  $struct_{ct}$ , or  $scalar_{ct}$ . The *conf*

<sup>1</sup> We are currently investigating methods for formally representing construct types and constructs, however, such a formal description is not required for the purpose of this paper.



**Fig. 3.** Relational, RDF, and XTM data models described within the generic framework.

predicate specifies that a construct (shown as  $c1$ ) can conform to another construct (shown as  $c2$ ). The  $ct-inst$  predicate relates an instance (shown as  $d$ ) with a construct (shown as  $c$ ). Finally, the  $ct-inst$  predicate defines a conformance relationship between instances (shown as  $d1$  and  $d2$ ).

We introduce the  $comp$  predicate for describing composition constraints over model constructs. The predicate can take the following forms.

```

comp(c1, setof(c1, ..., cn)).
comp(c2, listof(c1, ..., cn)).
comp(c3, bagof(c1, ..., cn)).
comp(c4, structof(a1→c1, ..., an→cn)).
comp(c5, union(c1, ..., cn)).

```

As shown, each kind of  $comp$  formula uses an additional formula to specify composition constraints, namely,  $setof$ ,  $listof$ ,  $bagof$ , or  $structof$ . For example,  $c1$  is assumed to be defined as  $ct-inst(c1, set_{ct})$  whose instances are sets that can contain instances of constructs  $c1$ ,  $c2$ , ...,  $cn$  (i.e., the  $setof$  predicate defines the homogeneity of the construct's instances). We also permit the definition of convenience-constructs (as shown by  $c5$  above), which do not have corresponding  $ct-inst$  formulas and must be associated with a  $union$  predicate. Such constructs serve as a place-holder for the union of the corresponding constructs.

The last predicate we define is  $val$ , which gives the value of an instance. Each ground  $val$  formula must have an associated  $ct-inst$  formula (which is also ground) of the appropriate form, as shown below.

```

val(d1, set(v1, ..., vn)).
val(d2, list(v1, ..., vn)).
val(d3, bag(v1, ..., vn)).
val(d4, struct(a1=v1, ..., an=vn)).
val(d5, scalar(v)).

```

To illustrate the definition language, Figures 4, 5, 6, and 7 define (simplified versions of) the XML with DTD, Relational, Topic Map, and RDF data models, respectively. Note that we assume  $string$  is a default scalar construct for each model and that the juxtaposition of entries is for exposition only. Finally, Figure 8 shows example instances for the XML model.

```

ct-inst(dtd, set_{ct}).
ct-inst(attr-t, struct_{ct}).
ct-inst(subelem-t, struct_{ct}).
ct-inst(doc, struct_{ct}).
ct-inst(attr, struct_{ct}).
ct-inst(nestedelem, struct_{ct}).
conf(doc, dtd).
conf(attr, attr-t).

ct-inst(elem-t, struct_{ct}).
ct-inst(pCDATA, scalar_{ct}).
ct-inst(attr-tset, set_{ct}).
ct-inst(elem, struct_{ct}).
ct-inst(attrset, set_{ct}).
ct-inst(cdata, scalar_{ct}).
conf(elem, elem-t).

comp(dtd, setof(elem-t)).
comp(elem-t, structof(tag→string, attrtypes→attr-tset)).
comp(attr-t, structof(name→string)).
comp(subelem-t, structof(parent→elem-t, child→elem-t)).
comp(attr-tset, setof(attr-t)).
comp(doc, structof(root→elem)).
comp(elem, structof(tag→string, attrs→attrset)).
comp(attr, structof(name→string, val→cdata)).
comp(attrset, setof(attr)).
comp(node, union(elem, pCDATA)).
comp(nestedelem, structof(parent→elem, child→node)).

```

**Fig. 4.** Simplified description of the XML with DTD data model.

```

ct-inst(db, set_{ct}).
ct-inst(table, struct_{ct}).
ct-inst(tuple, set_{ct}).
ct-inst(field, struct_{ct}).
ct-inst(value, scalar_{ct}).
conf(tuple, scheme).

ct-inst(scheme, set_{ct}).
ct-inst(tuples, bag_{ct}).
ct-inst(fieldtype, struct_{ct}).
ct-inst(domaintype, struct_{ct}).
conf(table, scheme).
conf(value, domaintype).

comp(db, setof(scheme, table)).
comp(scheme, setof(fieldtype)).
comp(table, structof(name→string, rows→tuples)).
comp(tuples, bagof(tuple)).
comp(tuple, setof(field)).
comp(fieldtype, structof(name→string, dom→domaintype)).
comp(field, structof(name→string, val→value)).
comp(domaintype, structof(name→string)).

```

**Fig. 5.** Simplified description of the Relational data model.

### 3.3 Specifying Constraints

In addition to describing composition of basic structures, we also provide a mechanism for defining the constraints of a data model. Namely, we leverage the logic-based description language for constraints by allowing constraint expressions (i.e., rules). We denote constraints using the annotation  $constraint$ , e.g., the following rule defines a conformance constraint over the Relational model.

```

ct-inst(tm, setct).          ct-inst(association, structct).
ct-inst(topic, structct).    ct-inst(assoc-mem, structct).
ct-inst(topic-occ, structct). ct-inst(occurrence, structct).
ct-inst(member, structct).   ct-inst(mem-topic, structct).
ct-inst(resource, structct). conf(topic, topic).
conf(association, topic).    conf(occurrence, topic).

comp(tm, setof(topic, assoc)).
comp(topic, structof(name→string)).
comp(topic-occ, structof(top→topic, occ→occurrence)).
comp(occ-val, union(resource, string)).
comp(occurrence, structof(val→occ-val)).
comp(association, structof()).
comp(assoc-mem, structof(assoc→association, mem→member)).
comp(member, structof(role→string)).
comp(resource, structof(uri→string)).

```

Fig. 6. Simplified description of the Topic Map (XTM) data model.

```

ct-inst(rdfmodel, setct).    ct-inst(uriref, structct).
ct-inst(literal, scalarct). ct-inst(blank, structct).
ct-inst(alt, setct).        ct-inst(seq, listct).
ct-inst(bag, bagct).        ct-inst(triple, structct).
conf(resource, resource).

comp(rdfmodel, setof(literal, blank, triple, coll)).
comp(resource, union(uriref, blank, coll, triple)).
comp(node, union(resource, literal)).
comp(uriref, structof(uri→string)).
comp(blank, structof()).
comp(coll, union(set, bag, alt)).
comp(alt, setof(node)).
comp(seq, listof(node)).
comp(bag, bagof(node)).
comp(triple, structof(pred→uriref, subj→resource,
                      obj→node)).

```

Fig. 7. Simplified description of the RDF data model.

(a).

```

<!ELEMENT activity (profession | hobbyist | ...)*>
<!ELEMENT profession (surgeon | professor | ...)*>
<!ELEMENT professor (#PCDATA)>

```

(b).

```

<activity>
  <profession>
    <professor>Lois</professor>
  </profession>
</activity>

```

(c).

```

c-inst(t, dtd).          val(t, set(et1, et2, et3)).
c-inst(t1, elem-t).     val(t1, struct(name='activity')).
c-inst(t2, elem-t).     val(t2, struct(name='profession')).
c-inst(t3, elem-t).     val(t3, struct(name='professor')).
c-inst(s1, subelem-t).  val(s1, struct(parent=t1, child=t2)).
c-inst(s2, subelem-t).  val(s2, struct(parent=t2, child=t3)).
c-inst(d, doc).         val(d, struct(root=e1)).
c-inst(e1, elem).       val(e1, struct(tag='activity')).
c-inst(e2, elem).       val(e2, struct(tag='profession')).
c-inst(e3, elem).       val(e3, struct(tag='professor')).
c-inst(n1, nestedelem). val(n1, struct(parent=e1, child=e2)).
c-inst(n2, nestedelem). val(n2, struct(parent=e2, child=e3)).
c-inst(l, pcdat).       val(l, scalar('Lois')).
c-inst(n3, nestedelem). val(n3, struct(parent=e3, child=l)).
c-inst(d, t).           d-inst(e1, t1).
d-inst(e2, t2).         d-inst(e3, t3).

```

Fig. 8. An example (a) XML DTD, (b) conforming document, and (c) both represented in the description language.

```

constraint: d-inst(X, Y) :-
  c-inst(X, table) & c-inst(Y, scheme) &
  not(d-inst(X, Yp) & Y≠Yp) &
  not(d-inst(Xp, Y) & X≠Xp) &
  val(X, B) & member(V, B) & d-inst(V, Y).

```

The constraint is read as: “X can conform to Y if X is a Table, Y is a Scheme, X doesn’t conform to any other Schemes, Y doesn’t have any other conforming Tables, and every Row in X conforms to Y.” Note that specifying constraints in this way provides a powerful mechanism for describing complex data model descriptions.

## 4 Transforming Representation Schemes

In this section, we present a language for transforming representation schemes based on horn-clause logic rules, provide example mappings, and discuss higher-level transformations for simplifying the specification and implementation of complex transformations.

### 4.1 A Transformation Language

A transformation is applied to one or possibly more configurations in which the result is a new configuration. Here, we only consider the case where a single configuration is transformed. We define a transformation as a function  $\sigma : M \times C \rightarrow C'$  where  $M$  is a set of *mapping rules*,  $C$  is the source configuration, and  $C'$  is the new configuration called the “destination” of the transformation. The transformation function  $\sigma$  computes the fix-point of the mapping rules applied to the source and destination configurations.

Mapping rules can take the following form:

$$dp_1 \ \& \ \dots \ \& \ dp_n \ :- \ p_1 \ \& \ \dots \ \& \ p_m,$$

where both  $dp$  and  $p$  are (a) formulas (using the predicates of Section 3) with ground literals or variables as arguments or (b) the formula  $ID(I, id)$ . The  $ID$  predicate takes a list  $I$  of input constants and provides a constant (*i.e.*,  $id$ ) that serves as a unique identifier for the input list. The  $ID$  predicate is implemented as a skolem function.<sup>2</sup> The syntax we use for mapping rules is shorthand for rules having a single, outer formula in the head of the program clause (via the head formula):

$$head(dp_1, \dots, dp_n) \ :- \ p_1 \ \& \ p_2 \ \& \ \dots \ \& \ p_m.$$

Mapping rules are read in the normal way, *i.e.*, if each  $p_i$  for  $1 \leq i \leq n$  is true, then each  $dp_j$  for  $1 \leq j \leq m$  is true. Additionally, we permit two annotations on mapping rule formulas. The `add` annotation indicates that the formula be added to the destination configuration. The `dest` annotation matches formulas against those in the destination configuration (the source configuration is the default). We do not allow transformations to modify the source configuration.

<sup>2</sup> We use the abbreviation  $ID(v_1+v_2+\dots, id)$  to denote the list  $[v_1, v_2, \dots]$

```

% element types become schemes
add:c-inst(S,scheme) :-
  c-inst(ET,elem-t) & ID(ET,S).
% attribute types become field types
add:c-inst(FT,fieldtype) :-
  c-inst(AT,attr-t) & ID(AT,FT).
% elements become tuples
add:c-inst(T,tuple) :-
  c-inst(E,elem) & ID(E,T).
% attribute become fields
add:c-inst(F,field) :-
  c-inst(A,attr) & ID(A,F).
% cdata to values
add:c-inst(V2,value) :-
  c-inst(V1,cdat) & ID(V1,V2).
% sub element types to schemes
add:c-inst(S,scheme) :-
  c-inst(ST,subelem-t) & ID(ST,S).
% field types for sub element type schemes
add:c-inst(FT1,fieldtype) &
add:c-inst(FT2,fieldtype) :-
  c-inst(ST,subelem-t) &
  val(ST,struct(parent→P,child→C) &
  ID(ST+P,FT1) & ID(ST+C,FT2).
% nested elements become tuples
add:c-inst(T,tuple) :-
  c-inst(NE,nestedelem) & ID(NE,T).
% each scheme has a table
add:c-inst(T,table) :-
  dest:c-inst(S,scheme) & ID(S,T).
% tables conform to their schemes
add:d-inst(T,S) :-
  dest:d-inst(S,scheme) & ID(S,T).
% values conform to the cdatatype domain type
add:d-inst(V,cdatatype) :-
  dest:c-inst(V,value).
% tuples conform to their schemes
add:d-inst(TP,S) :-
  c-inst(E,elem) & ID(E,TP) & d-inst(E,ET) & ID(ET,S).

```

Fig. 9. Rules for a model-to-model transformation.

## 4.2 An Example Model-to-Model Mapping

Figures 9 and 10 show a portion of a simple model-to-model mapping between XML and the Relational model using their descriptions from Section 3. Figure 9 shows instance and conformance rules and Figure 10 shows basic structure conversions. This mapping converts XML element types to Relational schemes and converts the rest of the data accordingly. Specifically, tables contain elements with their associated attribute values and a new table is created for each subelement-type relationship, which is where nested elements are stored. Figure 11 shows the tables that would result from example XML data. Note that we choose to skolemize each identifier to assure uniqueness in the destination (*e.g.*, the destination may contain other values prior to the transformation).

While model-to-model transformations are useful, there are a number of other transformations that can be supported in this framework, including schema-to-schema, model-to-schema (also call “modeling the model” [3]), and various mixtures of each [4, 5].

## 4.3 Higher-Level Transformations

While mapping rules are expressed using a declarative language, they specify the low-level details of transformations. Such an approach is attractive since it enables extremely flexible transformations (*e.g.*, model-to-model, schema-to-schema, and arbitrary combinations), however, specifying mapping rules is not trivial. In the remainder of this section, we describe two approaches for higher-level rule specification (each of which builds on the current mapping rule

```

% create an empty scheme for each element type
add:val(S,set()) :-
  c-inst(ET,elem-t) & ID(ET,S).
% add a field type for each attribute type
add:val(FT,struct(name=N,dom=cdatatype)) :-
  c-inst(AT,attr-t) & val(AT,struct(name=N)) &
  ID(AT,FT).
% add a field type for element ids
add:val(FT,struct(name='id',dom=stringtype)) :-
  c-inst(ET,elemtype) & ID(ET+'id',FT).
% add field types to associated element type schemes
add:member(FT,FTSet) :-
  c-inst(ET,elem-t) &
  val(ET,struct(tagname=_,attrtypes=ATSet)) &
  val(ATSet,Set) & member(AT,Set) & ID(ET,S) &
  ID(AT,FT) & dest:val(S,FTSet).
% create an empty scheme for each sub-element type
add:val(S,set()) :-
  c-inst(ST,subelem-t) & ID(ET,S).
% create two field types for sub element type scheme
add:val(FT1,struct(name=PN,dom=stringtype)) &
add:val(FT2,struct(name=CN,dom=stringtype)) :-
  c-inst(ST,subelem-t) &
  val(ST,struct(parrenttype→P,childtype→C) &
  val(P,struct(tagname→PN,attrtypes→_) &
  val(C,struct(tagname→CN,attrtypes→_) &
  ID(P,FT1) & ID(C,FT2).
% create a table element type
add:c-inst(R,tuples) & add:val(R,bag()) &
add:val(T,struct(name=N,rows=R)) :-
  c-inst(ET,elem-t) &
  val(ET,struct(tagname→N,attrtypes→_) &
  ID(ET,S) & ID(S,T) & ID(T,R).
% create a field for each attribute
add:val(F,struct(name→N,val→V)) :-
  c-inst(A,attr) & val(A,struct(name→N,val→C) &
  ID(A,F) & ID(C,V).
% create a field for each element as an id
add:val(F,struct(name→'id',val→V)) :-
  c-inst(E,elem) & ID(E+'id',F) & tostring(F, V).
...

```

Fig. 10. Structural rules for a model-to-model transformation.

approach), which we believe can make writing rules simpler. We briefly describe *transformation patterns* for capturing and re-using common transformation conventions and discuss semi-automatic *rule derivation*, which is enabled through our generic framework.

We view a transformation pattern as a mapping rule abstraction, *i.e.*, a transformation pattern is a parameterized specification taking mapping rules as input to create a more complex transformation. In this way, patterns are analogous to higher-order functions like `map` and `fold` in functional programming languages. In general, a pattern can be used as a runtime operation where a set of arguments is provided to perform the concrete transformation, or, as a mapping rule generator, *i.e.*, given a set of arguments, the pattern is used to generate (one or more) concrete mapping rules.

```

<!ELEMENT professor (person)>
<!ATTLIST professor dept CDATA #REQUIRED>
<!ELEMENT person EMPTY>
<!ATTLIST person name CDATA #REQUIRED>

<professor dept='CSE'>
  <person name='Lois' />
</professor>

```

Professor Table		Person Table		Professor-Person Table	
id	dept	id	name	professor	person
e1	"CSE"	e2	"Lois"	e1	e2

Fig. 11. Results from the XML to Relational transformation.

For example, there are often many ways to transform between models allowing multiple levels of conformance to models permitting only a single level. One such convention is to flatten multiple levels of conformance, making the conformance implicit in the destination. A different convention is to represent source instance-of relationships as hierarchical relationships in the destination (again, conformance becomes implicit). Examples are shown in Figure 12 where the source is a Topic Map and the target a Relational database (for the flattening approach) and an XML document (for the hierarchy approach). The source has three levels of conformance: topic “Lois” is an instance of topic “Professor”; topic “Professor” is an instance of topic “Profession”; and topic “Profession” is an instance of topic “Activity”.<sup>3</sup> In the case of XML, the hierarchical relationship is modeled as nested elements, and for the Relational example, each table represents a path of the original conformance relationship.

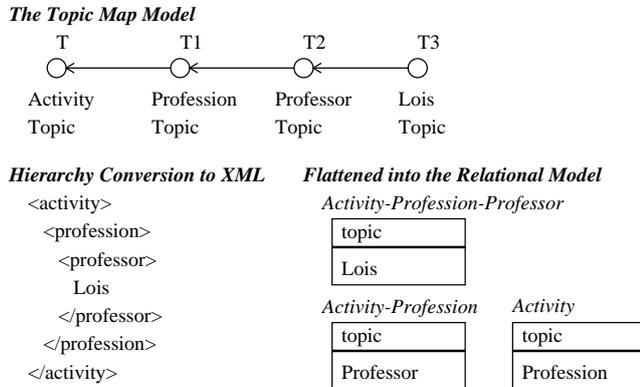


Fig. 12. An example of flattening and hierarchy conversion of the Topic Map model.

Figure 13 shows an example transformation pattern called `hierarchy-convert` for generically applying the hierarchy conversion. (The language shown for describing patterns is for exposition only; we are currently developing a pattern language for transformations.) The pattern takes five arguments: the construct instance `T` serving as the top item of the multi-level conformance relationship (e.g., “Activity” in Figure 12); a mapping rule  $M_1$  to map source items to target items (e.g., topics to XML elements); a special mapping rule  $M_2$  for converting items at the leaf level (e.g., instead of mapping topics to elements, we would map topics to XML element content); and mapping rules  $A_1$  and  $A_2$  for attaching the hierarchical relationships in the destination.

The `hierarchy-convert` pattern is broken into two sub-rules. The first rule identifies items above the leaf conformance items (e.g., the “Activity” and “Profession” topics) and converts them to target items (using  $M_1$ ). The second rule is similar, but for leaf nodes. In our example, the second rule would convert the topic “Lois” into XML content (using  $M_2$ ), and attach it to the “Professor” element (using  $A_2$ ).

We see a number of potential patterns as useful, e.g., a simple class of patterns are those that convert between construct

<sup>3</sup> To see this is correctly modeled, notice that “Lois” is an instance of a “Professor,” but not an instance of a “Profession.”

```

pattern:hierarchy-convert(T,M1,M2,A1,A2) =
% convert inter-nodes T and T1
map T to S and T1 to S1 using rule M1 &
attach S1 to S using rule A1 :-
d-inst(T3,T2) & d-inst(T1,T) &
c-inst(T,C) & c-inst(T1,C) & c-inst(T2,C).
% convert last node
map T2 to S2 using rule M1 &
map T3 to S3 using rule M2 &
attach S3 to S2 using rule A2 :-
d-inst(T3,T2) & not d-inst(T4,T3) &
c-inst(T2,C) & c-inst(T3,C) & c-inst(T4,C).
    
```

Fig. 13. Example of the conformance flattening pattern.

types. That is, patterns for converting lists to sets and bags, flattening nested sets, merging sets, lists, bags, and structs, and converting a group of structs into sets, and so on.

Patterns help make it easier to write rules by enabling the reuse of commonly used transformations. Another approach for simplifying rule writing is semi-automatic rule derivation. The goal is to generate the structural mapping rules (like those in Figure 10) whenever possible from the simpler construct level rules (such as those in Figure 9) by leveraging the composition constraints of constructs.

For example, consider the excerpt of XML and Relational models shown in Figure 14. The following mapping rule specifies a construct transformation between XML element types and Relational schemes.

```

c-inst(ET,scheme) :- c-inst(ET,elem-t).
    
```

Based on the description and constraints of the source model, we can conclude the following: (1) every `scheme` must have exactly one `table` (and vice versa), (2) every `table` is a bag of `tuple`'s, (3) every `tuple` must conform to its `table`'s conforming `scheme`, and (4) more than one `elem` can conform to an `elem-t`.

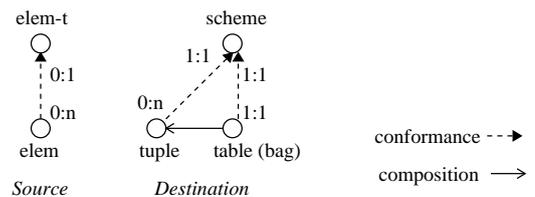


Fig. 14. Source and destination constructs for rule derivation.

From the above mapping rule requesting each `elem-t` be converted to a `scheme` and the above constraints, it is reasonable to consider the following mappings.

1. Each `elem` should be mapped to a `tuple` since many `elem`'s can be associated with each `elem-t`, and it is only through `tuple`'s that multiple items can be associated with a given `table`.
2. A `table` should be created for each converted `elem-t` (i.e., each new `scheme`).
3. The corresponding `table` (for the new `scheme`) should contain the converted `tuple` items.

Note that to generate these rules, we can employ structural transformation patterns, e.g., to convert the source `d-inst`

relations into a bag (table) of the appropriate tuples, and so on. Although a simple example, the derived rules are equivalent to those used to map element types to schemes in the transformation of Section 4.2. Thus, combining patterns and constraint-based rule derivations can ease the task of rule specification, both of which are enabled by the uniform framework. We intend to investigate both patterns and rule derivation further.

## 5 Related Work

In this paper, we extend our earlier work [4, 5], which introduced the use of a uniform, RDF-based representation scheme allowing model, schema, and instance information to be represented in triples. First, the metamodel is significantly more complex. It also has a relatively complete set of structural primitives, including most database modeling constructs: *struct* (much like an object in object-oriented models) and explicit collection primitives *set*, *bag*, and *list*. More importantly, these basic structures can be freely composed as appropriate for the data model being described using the constraint facility. We also contribute a clear separation of basic structures (*i.e.*, construct types) from instance-of and conformance relationships—this separation differs from RDF and Topic Map models each having only a single version. Thus, the choice of construct types (*i.e.*, metamodel) is extensible and the architecture can be exploited to provide high-level transformation rules (*i.e.*, patterns and constraint-based rule derivation).

The metamodel architecture we propose is the only approach we know of that explicitly models conformance. Existing metamodel approaches [6–11] primarily focus on solving database interoperability issues and assume (and exploit) schema-first data models. Typically, transformation is defined at the data model level or at the schema level (but not both). The transformation approach we propose enables mappings at differing levels of abstraction, *e.g.*, model-to-model, model-to-schema, and combinations of these. In addition, we permit partial mappings, *i.e.*, we do not require complete (or one-to-one) transformations.

The use of logic-based languages for transformation has been successfully applied [10, 12], *e.g.*, WOL [12] is a schema transformation language that uses first-order logic rules for both schema transformations and specifying (schema) constraints. Similar to WOL, we use a logic-based approach for specifying constraints.

Finally, a number of recent model-to-model and model-to-schema mappings (without a formal description of the transformation) have been defined between XML and Relational models [13] and between Topic Maps and RDF [3, 14]. We believe using a formal transformation language can benefit both defining such mappings, and implementing them.

## 6 Summary and Future Work

The three-level architecture we propose has the following advantages. It allows the data model designer to explicitly define when and how one construct conforms to another (through conformance links). It can be used to describe models that require schema, don't require schema, permit multi-

ple schemas, allow multiple levels of schema-instance connections, and include any or all of these possibilities. It introduces four explicit relationships, *ct-inst*, *conformance*, *c-inst*, and *d-inst*, which can be easily exploited in a transformation language as well as by other tools such as a query language or a browser. Finally, it provides orthogonal, complementary specification of conformance and instance-of relationships versus composition and constraints for describing actual data structures that exist in the information representation.

We believe patterns and constraint-based transformation offer powerful mechanisms for specifying mapping rules and intend to further investigate their use. We are also exploring “in-place transformation” for performance optimization, *i.e.*, transforming information without first converting it to an intermediate representation. Thus, mappings are implemented by calls to the associated source and destination native technology. Such native rule generation is made possible by the specification of the data model using the framework we have presented.

## References

1. Marshall, C., III, F.S., Coombs, J.: Viki: Spatial hypertext supporting emergent structure. In: European Conf. on Hypertext Technology (ECHT '94), ACM (1994) 13–23
2. Nanard, J., Nanard, M.: Should anchors be typed too? An experiment with macweb. In: Proc. of Hypertext. (1993) 51–62
3. Moore, G.: RDF and TopicMaps: An exercise in convergence. In: XML Europe. (2001)
4. Bowers, S., Delcambre, L.: Representing and transforming model-based information. In: Proc. of the First Inter. Workshop on the Semantic Web. (2000)
5. Bowers, S., Delcambre, L.: A generic representation for exploiting model-based information. In: ETAI Journal. Volume 6. (2001)
6. Atzeni, P., Torlone, R.: Management of multiple models in an extensible database design tool. In: EDBT. Volume 1057 of Lecture Notes in Computer Science. (1996) 79–95
7. Atzeni, P., Torlone, R.: A unified framework for data translation over the web. In: Proc. of WISE. IEEE Computer Society Press (2001)
8. Barsalou, T., Gangopadhyay, D.: M(DM): An open framework for interoperation of multimodel multidatabase systems. In: Proc. of ICDE. (1992) 218–227
9. Christophides, V., Cluet, S., Siméon, J.: On wrapping query languages and efficient XML integration. In: Proc. of SIGMOD. (2000) 141–152
10. McBrien, P., Poulouvasilis, A.: A uniform approach to inter-model transformations. In: Proc. of CAiSE'99. Volume 1626 of Lecture Notes in Computer Science. (1999) 333–348
11. Papazoglou, M., Russell, N.: A semantic meta-modeling approach to schema transformation. In: Proc. of CIKM, ACM (1995) 113–121
12. Davidson, S., Kosky, A.: WOL: A language for database transformations and constraints. In: Proc. of ICDE. (1997) 55–65
13. Bohannon, S., Freire, J., Roy, P., Siméon, J.: From xml schema to relations: A cost-based approach to xml storage. In: ICDE. (2002)
14. Lacker, M., Decker, S.: On the integration of Topic Maps and RDF. In: Proc. of the 1st International Semantic Web Working Symposium (SWWS '01). (2001)