

ATLANTIC INTERNATIONAL UNIVERSITY

A Model Theory for Generic Schema Management Models

Kinan M Al Haffar

UM4699SIT10550D

07/08/2007

Abstract

The core of a model theory for generic schema management is developed. This theory has two distinctive features: it applies to a variety of categories of schemas, and it applies to transformations of both the schema structure and its integrity constraints. A subtle problem of schema integration is considered in its general form, not bound to any particular category of schemas. The proposed solution, as well as the overall theory, is based entirely on schema morphisms that carry both structural and semantic properties. Duality results that apply to the two levels (i.e., the schema and the data levels) are established. These results lead to the main contribution of this paper: a formal schema and data management framework for generic schema management. Implications of this theory are established that apply to integrity problems in schema integration. The theory is illustrated by a particular category of schemas with object-oriented features along with typical database integrity constraints.

1 Introduction

This paper presents the core results of a model theory for generic schema management, by which we mean schema and database transformation capabilities that are independent of a particular data model. Such transformations require major database programming tasks, such as integrating source schemas when building a data warehouse or integrating different user views into an overall database schema. In spite of nontrivial typing issues created by such transformations, database programming and other relevant paradigms have been primarily suited to dealing with structural aspects of those transformations. A major challenge is in properly addressing semantics: the integrity constraints associated with database schemas.

A second major challenge is in developing such a model theory that is applicable to a variety of data models, such as the relational, object-oriented, and XML models [23]. This is challenging because schemas and their underlying databases are very different in these three major categories of models as are languages and their underlying logics used for expressing the integrity constraints.

On the pragmatic side, generic schema management operations and tools have been considered in [4], and more specific system implications in [5, 20]. These papers argue that many difficult and expensive database programming problems involve the manipulation of mappings between schemas. Examples are populating a data warehouse from data sources, exposing a set of data sources as an integrated schema, generating a web site wrapper, generating an object-oriented wrapper for relational data, and mapping an XML schema to a relational schema. Despite many commonalities of these schema management problems, tools and languages are typically engineered for only one problem area. A more attractive approach would be to build generic schema management tools and languages that would apply to all of these problems, with only some customization required for the data model and problem at hand. The formal open problem is to find a suitable model theory that would be able to handle this generality.

Our proposed model theory has a categorical flavor, manifested in the use of arrows that appear at two levels. At the meta level the arrows represent schema transformations. For example, an arrow could map a data source schema to a data warehouse schema, or two arrows could map each of two data source schemas into a mediated schema. These transformations are defined as schema morphisms that map the structural properties and integrity constraints of a schema. At the data level the arrows are data transformations specified as database morphisms. Database morphisms map the actual data sets in a manner that is compatible with the operations available on those sets and that preserves the schema's integrity constraints.

There are several implications of this arrow-theoretic approach. The first is that it leads to a very general view of the schema integration problem expressed entirely in terms of arrows. The generality is accomplished by using a particular categorical construction which applies to a variety of categories of schemas [10, 12, 17]. The specific nature of arrows used in this construction is determined by the category

of schemas that defines the kind of structural and integrity transformations that the arrows actually represent.

The second implication is in the formal results that relate properties of arrows at the meta level and the data level. These results reveal a subtle duality manifested in the reversal of the corresponding arrows between the two levels. They also provide guidance for how to define schema transformations appropriately, in particular, so that the arrows preserve integrity constraints.

The third, most important implication is a formal framework for generic schema and data management which applies to a variety of data models. This is really the main contribution of this paper. This formal framework includes and relates the two levels (schema and data transformations) and captures both structural properties and integrity constraints. A key component of this theory is a strong integrity requirement on the permissible structural schema transformations so that they preserve the integrity constraints. This model theory relies on earlier results on a general model theory for a variety of programming paradigms and of associated logic paradigms [9, 10].

We apply the proposed formal theory to two situations. First, we apply it to the schema integration problem where we prove results that are generic across data models and include a proper treatment of the integrity constraints. Second, we use it as a pattern for defining a data model, one based on abstract data types, which includes object-oriented features and typical database integrity constraints.

The paper is organized as follows. Section 2 starts with basic definitions, followed by a running example of a particular category of object-oriented schemas. The basic categorical definitions are given in Section 3, particularly the use of morphisms to represent schema mappings. Sections 4 through 6 are the core of the paper. Section 4 shows how to express the schema integration problem using morphisms. Section 5 introduces the notion of a generic schema transformation framework, specifies the condition for such a framework to preserve database integrity and proves the implications on schema integration. Section 6 shows how the proposed model theory applies to the category of schemas of the running example. Section 7 discusses related work and concludes.

2 Schemas

A database schema consists of two components: a schema signature and the associated integrity constraints. A schema signature specifies structural and operational features of a database schema. Its typical components are signatures for data types and their operations and signatures for *database sets* (relations, collections, etc.), illustrated in Example 1. Note that type signatures for collections determine the signatures for operations on collections. Signatures for integrity constraints are logical expressions whose form is determined by the choice of a particular logic and syntax of the constraint language. A feature of the definitions below is that they are sufficiently general to apply to a variety of schema signatures and associated constraint languages.

Definition 1 (*Schemas*) A database schema Sch is a pair $Sch = (Sig, E)$ where Sig is a schema signature and E is a set of integrity constraints expressed as sentences (formulae of a chosen logic with all variables quantified).

This paper uses examples based on a category of schemas of an object-oriented style. Schemas of this category have user-defined abstract data types specified as Java interfaces. A schema signature consists of the specification of those types and of collection objects which represent the actual database. The integrity constraints are specified in Horn clause logic. We use only the most typical database constraints: those expressing key dependencies and referential integrity constraints.

Example 1 (*Sample schema*)

```
schema Publishers { interface
Publisher { String publisherIdO ;
    String name();
    String locationO ;
    Set<Publication> publications();
}
interface Publication {
String publicationIdO ;
    String title();
    int year();
    Publisher publisherO ;
```

```

Set<String> keywords ( ) ;
}
Collection<Publisher>    dbPublishers;
Collection<Publication> dbPubs; Constraints :
Publication X,Y;  Publisher Z,W;
dbPubs.member(X)      :- dbPublishers.member (Z) , Z .publications.member(X);
dbPublishers.member(Z) :- dbPubs.member(X) , X.publisher.equals(Z);
Z.equals(W)  :- dbPublishers.member(Z) , dbPublishers.member(W) ,
                Z.publisherIdO .equals (W.publisherIdO ) ; X.equals(Y)  :- dbPubs .member (X) ,
dbPubs.member(Y) , X.publicationIdO . equals (Y.publicationIdO ) ; }

```

As Java interfaces lack any kind of general-purpose constraint specifications, such constraints are omitted from our examples. However, this omission is by no means a limitation of our approach (see [1, 2]). Note that equals is a method of the Java root class **Object** which is intended to be overridden to provide a meaning of equality specific to a particular data type. If this is the standard notion of equality then the logic paradigm becomes Horn clause logic with equality as in [11].

A database is a model for a schema. Given a schema signature Sig , the collection of all databases that *conform to Sig (implement Sig)* is denoted by $Db(Sig)$. A database d that implements Sig would thus have to implement the signatures of data types in Sig as sets and the operation signatures as functions. Signatures for database sets would also have to be interpreted as sets, bags etc.

If a database d that conforms to a schema signature Sig satisfies the integrity constraints E , we say that d is consistent with respect to the schema (Sig,E) . This is expressed by the satisfaction relation, denoted \models , between databases and the sets of sentences (i.e., integrity constraints) that they satisfy.

Definition 2 (*Database consistency*) A database d is consistent with respect to the schema (Sig,E) iff d belongs to $Db(Sig)$ and $d \models e$ for all $e \in E$.

3 Schema Morphisms

In this approach the schema signatures for a given data model are required to constitute a category, denoted **Sig**. The same applies to the schemas. Schema transformations within a particular category of schemas are viewed as morphisms of that category. This approach allows us to talk about schemas without specifying their data model (i.e. category). Formally, a *category C* consists of the following [17]:

- A collection of objects and a collection of arrows (morphisms).
- Each arrow has its domain object and its codomain object.
- If $f : X \longrightarrow Y$ and $g : Y \longrightarrow Z$ are arrows of C , then f and g are composed into an arrow $gf : X \longrightarrow Z$.
- The composition of arrows is associative, i.e., if $f : X \longrightarrow Y$, $g : Y \longrightarrow Z$ and $h : Z \longrightarrow W$ are arrows of C , then $h(gf) = (hg)f$.
- Each object X of C is equipped with the identity arrow 1_x with the property that for any arrow $f : X \longrightarrow Y$, $1_y f = f$ and $f 1_x = f$.

The collection of all databases $Db(Sig)$ conforming to a given schema signature Sig is required to constitute a category, with database morphisms that satisfy the above categorical requirements.

Schema morphisms are defined as mappings of schema signatures that preserve the integrity constraints.

Definition 3 (*Schema morphisms*) A morphism of schemas $Sch1 = (Sig_1, E_1)$ and $Sch2 = (Sig_2, E_2)$ consists of a morphism $\langle j \rangle : Sig_1 \longrightarrow Sig_2$ of schema signatures which extends to a mapping of constraints, such that for all databases d in $Db(Sig_2)$, if $d \models e_2$ for all $e_2 \in E_2$ then $d \models \langle j \rangle(e_1)$ for all $e_1 \in E_1$.

An equivalent condition is that $\langle j \rangle(e_1)$ is in the closure of E_2 for all $e_1 \in E_1$. A particular case is $\langle j \rangle(e_1) \in E_2$. Definition 3 implies that schemas and their morphisms constitute a category, which we denote by **Sch**. The last condition in Definition 3 differentiates this work from many others. It requires that the integrity constraints of the source schema are transformed into constraints that are consistent with those of the target. This is expressed in a model-oriented fashion: a database that satisfies the constraints of the target schema also satisfies the transformed constraints of the source schema (as they appear in the target).

In Definition 1, distinguishing the notions of schema signature and schema (i.e., a schema signature plus its integrity constraints) leads to two notions of schema equivalence. The first one is just structural, based on schema signature morphisms. The second is semantic, requiring structural equivalence plus the semantic equivalence expressed in terms of integrity constraints.

Definition 4 (*Schema equivalence*) Let $Schi$ and $Schz$ be two schemas

- $Schi$ and $Schz$ are structurally equivalent if there exists a pair of schema signature morphisms $f : Sigi \rightarrow Sig_2$ and $g : Sig_2 \rightarrow Sigi$ such that $fg = Id_{Sig_2}$ and $gf = Id_{Sigi}$ ■
- $Schi$ and $Schz$ are equivalent if the above condition is satisfied for schema morphisms f and g .

The category of all schemas for a given data model is required to include the initial schema $Scho$ that contains basic features implicitly available in any other schema (for example, predefined standard types such as Boolean, integer and string, and their associated constraints). $Scho$ typically does not contain signatures for database sets but does contain type signatures for the required collection types. Any other schema $Schx$ implicitly extends $Scho$ by a unique schema morphism $Scho \rightarrow Schx$. The uniqueness requirement means that there is one standard way of incorporating $Scho$ into $Schx$ which makes $Scho$ a subschema of $Schx$.

The notion of a subschema $Schi$ of $Schx$ may be expressed by the categorical requirement that a monic arrow (schema morphism) $m : Schi \rightarrow Schx$ exists. This means that given schema morphisms $f : Schx \rightarrow Schi$ and $g : Schx \rightarrow Schi$, $mf = mg$ implies $f = g$ [17]. In familiar cases, monic arrows are injections.

4 Schema Integration

A particularly important problem in schema management is schema integration [6, 21, 22]. Here, we are given two (or more) schemas and are asked to produce an integrated schema that can represent the information content of the given schemas. In one version of the problem, which we treat here, the integrated schema is populated with data and the given schemas are defined as views of the integrated schema.

All published schema integration results we know of are tied to a particular data model (i.e., to a particular category of schemas). By contrast, our approach applies to different categories of schemas. Moreover, it addresses the subtle problem of merging the integrity constraints of two schemas. These two distinctive features are accomplished by expressing the idea of schema integration in terms of morphisms of schemas. This way both structural and semantic conditions are taken into account.

Definition 5 (*Schema integration*) An integration $Schi_2$ of schemas $Schi$ and $Schi$ is defined by the following commutative diagram of schema morphisms:

$$\begin{array}{ccc} Sch_m & \longrightarrow & Sch_i \\ hi & & P\{ \\ Sch_i & \longrightarrow & Sch_2 \end{array}$$

Two schemas are integrated over their matching part Sch_m . Two schemas always have a matching part: the initial schema $Scho$. The matching part Sch_m is typically a subschema of both $Schi$ and $Schi$. The commutativity of the above diagram asserts that the two composite schema morphisms $pcf \circ i$ and qfa are identical. This means that Sch_m appears the same in Sch_n whichever path (via $Schi$ or $Schi$) is taken.

Example 2 (*Schema integration*) We show the integration $Schi_2$ of $Schi$ (in Example 1) with another schema, $Schz$, such that $Schi$ and $Schi$ can be defined as views of Sch_n . First, we define $Schi$ as follows:

```
schema Authors {
interface Author {
String authorIdO ;
    String name();
    Date dateOfBirthO ;
    Collection<Book> books(); }
interface Book { String
publicationlDO ;
```

```
String title();
int year();
String publisherO ;
Set<String> keywords();
float price();
Collection<Author> authors();
}
Collection<Author> dbAuthors;
Collection<Book> dbBooks;

Constraints:
Author X,Y; Book Z,W;
dbBooks.member(Z) :- dbAuthors.member(X), X.books.member(Z);
dbAuthors.member(X) :- dbBooks.member(Z), Z.authors.member(X);
X.equals(Y) :- dbAuthors.member(X), dbAuthors.member(Y), X.authorldO .equals(Y.authorldQ );
Z.equals(W) :- dbBooks.member(Z), dbBooks.member(W), Z.ISBN().equals(W.ISBN())
}
```

A schema $Schn$ that integrates $Schi$ and $Schz$ is given below:

```
schema Publications {
import Publishers.Publisher;
import Publishers.Publication;
import Authors.Author;
interface Book extends Publication // integrates Book and Publication in Sch12
{ float price();
Collection<Author> authors(); }
Collection<Publisher> dbPublishers;
Collection<Publication> dbPubs;
Collection<Author> dbAuthors;
Collection<Book> dbBooks; Constraints :
import Publishers.Constraints; import
Authors.Constraints; Publication X;
dbPubs.member(X) :- dbBooks.member((Book)X) // an additional constraint beyond Sch1 and Sch2 }
```

Among all integrations $Schn$ of schemas S_1 and S_2 , the schema-join of $Schi$ and Sch_2 denoted $Sch_1 * Sch_2$, if it exists, has a distinctive property: it represents the minimal integration of $Schi$ and Sch_2 . This notion is specified below entirely in terms of schema morphisms by a categorical construction called a *pushout* [17].

Definition 6 (Schema join) *Schema-join $Sch_1 * Sch_2$ of schemas $Schi$ and Sch_2 is defined by the following commutative diagram of schema morphisms:*

$$\begin{array}{ccc}
 Sch_m & \longrightarrow & Sch_i \\
 \mathbf{hi} & & \mathbf{hf} \\
 Sch_i & \longrightarrow & Sch_i * Sch_i
 \end{array}$$

with the following property: Given any integration $Schn$ of schemas $Schi$ and Sch_2 as defined in Definition 5 above, there exists a unique schema morphism $\langle f \rangle : Sch_i * Sch_i \longrightarrow Sch_n$ such that $\langle f \rangle \circ k = q$ and $\langle f \rangle \circ h = p$.

A specific construction of schema-join is presented in Section 6. The fact that schema integration and schema-join are expressed entirely in terms of arrows representing schema morphisms has two distinctive implications: (i) both structural and database integrity properties are integrated, and (ii) the notion of schema integration is data model independent. Specific notions of schema integration are obtained by choosing a particular category of schemas which includes its schema morphisms.

We believe that the notion of least upper bound in the schema lattice of [7] and the schema + operator of [19] can be modeled precisely as schema join. Our category-theoretic treatment generalizes [7] by considering the instance level, not just the schema level and generalizes both models by making it applicable to a variety of schema categories, including such features as methods and constraints.

5 Generic Schema Transformation Framework

The categorical approach developed so far leads to a very general notion of a data model. This model has two levels. The meta level consists of database schemas and their transformations expressed as schema morphisms. This transformation-based view is quite different from the standard notions of a data model. At the instance level these transformations operate on databases that conform to schemas at the meta level.

The notion of database integrity plays a fundamental role in this formal framework. The acceptable transformations at both levels are required to satisfy the integrity constraints. This integrity requirement is expressed as a condition that involves acceptable schema signature transformations and the satisfaction relation between databases and the integrity constraints.

The framework is proposed as a formal pattern for defining data models such that schema management and the associated database transformations have well-defined formal meanings. As such, it offers interesting observations on the relationships between the two levels which are sometimes quite different from the usual views of schema and data transformations. An example is the reversal of the directions of transformations at the two levels. Schema management operations such as schema integration have particularly desirable semantic properties when the conditions for this formal framework are satisfied.

The core of this model theory requires just one more categorical notion: a morphism of categories. Given categories C and B , a *functor* $F : C \rightarrow B$ consists of two related functions with the following properties [17]:

- The object function which assigns to each object X of C an object FX of B .
- The arrow function which assigns to each arrow $h : X \rightarrow Y$ of C an arrow $Fh : FX \rightarrow FY$ of B .
- $F(lx) = lFX$ and $F(gh) = F(g)F(h)$, the latter whenever the composite gh is defined.

Two functors play a crucial role in this theory. The first one is $Sen : \mathbf{Sign} \rightarrow \mathbf{Set}$ where \mathbf{Set} denotes the category of sets. If Sig is a schema signature, then $Sen(Sig)$ is the set of all well-formed sentences (integrity constraints) over Sig . Many constraint languages, hence many Sen functors, are possible for a given data model (i.e., which implies a choice of \mathbf{Sign}). Sen is determined by the choice of logic that specifies the syntax of sentences of the constraint language. This syntax is defined starting with the features of a schema signature. The schema signature typically determines the terms of the constraint language, and the logic determines the formulae based on those terms. This is why Sen maps a schema signature into a set of sentences over that signature. Sen also maps a schema signature morphism to a function that transforms the sets of sentences.

The second functor is $Db : \mathbf{Sign} \rightarrow \mathbf{Cat}^{op}$. For each signature Sig , $Db(Sig)$ is the category of Sig databases, together with their morphisms. These database morphisms represent data transformations, which correspond to the schema transformations represented by arrows in \mathbf{Sign} . \mathbf{Cat} denotes the category of categories. Objects of \mathbf{Cat} are categories and arrows of \mathbf{Cat} are functors. \mathbf{Cat}^{op} differs from \mathbf{Cat} only to the extent that the direction of its arrows is reversed. This reversal of the direction of arrows that happens going from schemas to their databases is one of the subtle and characteristic features of this model theory.

Recall that databases in the category $Db(Sig)$ are not necessarily consistent with respect to a set of integrity constraints E . The notion of database integrity is captured by the satisfaction relation \models between databases in $Db(Sig)$ and sets of sentences over Sig .

We now have the machinery to define the notion of data model described in the beginning of this section. To emphasize the transformation-based view relative to classical data models, we give it a new name.

Definition 7 (*Schema transformation framework*) A schema transformation framework consists of:

- A category of schema signatures \mathbf{Sign} equipped with the initial object. This category consists of objects representing schema signatures together with their morphisms.
- A functor $Sen : \mathbf{Sign} \rightarrow \mathbf{Set}$. $Sen(Sig)$ is a set of sentences over the schema signature Sig .
- A functor $Db : \mathbf{Sign} \rightarrow \mathbf{Cat}^{op}$. For each signature Sig , $Db(Sig)$ is the category of Sig databases, together with their morphisms.
- For each signature Sig , a relation $\models_{sig} Q \mid Db(Sig) \times Sen(Sig)$ called the satisfaction relation. $\setminus Db(Sig) \setminus$ denotes the set of objects of the category $Db(Sig)$.
- For each schema signature morphism $(j) : Sig A \rightarrow Sig B$, the following Integrity Requirement holds for each $Sig B$ database dB and each sentence $e \in Sen(Sig A)$:

$$d_B \models_{sig B} Sen(j)(e) \text{ iff } Db(j)(d_B) \models_{sig A} e.$$

The above definition is based on [9]. Its relationships are represented by the following diagram:

$$\begin{array}{ccc}
 Db(SigA) & \longrightarrow & Sen(SigA) \\
 Db^{\wedge} & & | Sen(0) \\
 Db(SigB) & \longrightarrow & Sen(SigB)
 \end{array}$$

Note the reversal of the direction of the arrow $Db(\langle j \rangle)$ relative to $Sen(\langle j \rangle)$. $Sen(\langle f \rangle)$ maps each constraint in $Sen(SigA)$ to a constraint in $Sen(SigB)$. By contrast, $Db(\langle f \rangle)$ maps each database in $Db(SigB)$ to a database in $Db(SigA)$. If we think of $\langle f \rangle$ as a mapping from a logical database schema $SigA$ into a physical schema $SigB$, then $Db(\langle f \rangle)$ tells how to materialize a view in $Db(SigA)$ from a database in $Db(Sigs)$.

To see why this reversal of arrows happens, consider an injective schema morphism $\langle f \rangle : Sch \rightarrow Sch^{\wedge}$ which makes Sch a subschema of Sch^{\wedge} . $Schi$ just extends the signatures and the constraints of Sch (we assume a monotonic logic such as Horn clause logic). A database that is consistent with respect to $Schi$ is also (by projection) consistent with Sch . The other way around does not hold.

The *Integrity Requirement* puts a very strong semantic restriction on the permissible schema signature transformations (schema morphisms). It only allows ones that preserve the validity of constraints. That is, suppose $\langle f \rangle$ maps constraint BA to es (more precisely, $es = (Db(\langle j \rangle))(eA)$ for $BA \in Sen(SigA)$). Then the *Integrity Requirement* says that es is valid in a database $(B$ iff $\&A$ is valid in the $SigA$ database that corresponds to ds (i.e., in $Db(\langle j \rangle)(d,B)$).

The *Integrity Requirement* applies directly to the problem of schema integration in Definition 6. It ensures that each valid database of the integrated schema corresponds to valid databases of the schemas to be integrated. This point is made precise in the following theorem.

Theorem 1 (Materializing subschemas) Suppose Sch and Sch^{\wedge} are schemas that are integrated into $Schi_2$ relative to some schema transformation framework. Given a consistent database d in $Db(Sch_2)$, it is possible to construct databases d_1 and d_2 that are consistent relative to schema Sch and schema $Schi$ respectively.

Proof We have schema morphisms $\langle f \rangle_1 : Sch \rightarrow Sch_2$ and $\langle f \rangle_2 : Sch \rightarrow Sch_2$ where $Sch = (Sig, E)$ and $Sch_2 = (Sig_2, E_2)$. We construct database d_1 as $Db(\langle f \rangle_1)(d)$ and database d_2 as $Db(\langle f \rangle_2)(d)$. Since d is a consistent database and $\langle f \rangle_1$ and $\langle f \rangle_2$ are schema morphisms, Definition 3 implies:

- $d \models_{sj_1} (Sen(\langle f \rangle_1)(e))$ for all $e \in E_1$
- $d \models_{sj_2} (Sen(\langle f \rangle_2)(e))$ for all $e \in E_2$.

We can now complete the proof by applying the Integrity Requirement to the above two lines, yielding:

- $Db(\langle f \rangle_1)(d) \models_{si_1} e$ for all $e \in E_1$
- $Db(\langle f \rangle_2)(d) \models_{si_2} e$ for all $e \in E_2$.

In essence, databases conforming to Sch and Sch_2 are materialized views of $Schi_2$. Thus, the above theorem relates the consistency of databases conforming to the integrated schema to those of the materialized views. This is an unusual perspective in that views typically do not have integrity constraints.

Note that in order to make this proof possible both the reversal of arrows and the Integrity Requirement are essential.

Corollary 1 (Schema-joins) Let $Schi$ be an integration of schemas Sch and Sch_2 relative to a schema transformation framework. Let d be a consistent database in the category $Db(Sig_2)$. If $Sch * Sch_2 = (Sig_2, E_2)$ exists, then there is a canonical way to construct a consistent database d^* in the category $Db(Sig_2)$.

Proof By Definition 6, there is a unique $\langle f \rangle : Sch * Sch_2 \rightarrow Sch_2$. d^* is constructed as $Db(\langle f \rangle)(d)$. Since d is consistent, by Theorem 1 so is d^* .

Corollary 1 says that for any consistent database of an integrated schema, there exists a corresponding (minimal) consistent database of the schema-join of the two schemas that were integrated. The canonical arrow $Db(\langle f \rangle)$ is a recipe for constructing that (minimal) consistent database.

6 An Application of the Schema Transformation Framework

One role of the generic schema transformation framework is to serve as a formal generic pattern for defining new data models, with well defined meanings for schema and data transformations. This section gives a detailed description of how one constructs such a data model: a category of schemas and their databases for Examples 1 and 2, which we call **Objc** (OO with Constraints).

A schema of **Objc** consists of a collection of sorts (type names) some of which are predefined in the initial schema *Scho* and which in particular must include the sort Boolean along with the standard axioms. The others are abstract data types defined in the schema itself, each of which is a set of method signatures specified as a Java interface. The is-a (inheritance) relationship thus amounts to the subset relation which agrees with the rules of the Java type system. A schema contains collection objects which are the actual database sets. To specify their type, a parametric *Collection* $\langle T \rangle$ type is used and instantiated with the required element type selected among the abstract data types defined in the schema.

Definition 8 (Objc schema signatures) A schema signature consists of:

- A finite set of sorts S , which includes Boolean.
- A finite set of interfaces A (abstract data types) such that $A \subseteq S$.
- An interface A_k is a set of method signatures of the form $C \ m(C \setminus x_i, C_i \ x_2, \dots, C_n \ x_n)$ where $C \in S$ and $C_i \in S$ for $i = 1, 2, \dots, n$. (m is the method name, C is the return type, and $C \setminus$ is the type of parameter x_i .)
- If the interface $A \wedge$ extends the interface $A \setminus$ (inheritance) then $A \setminus \subseteq A \wedge$.
- A finite set of collection signatures $\{Collection \langle A_j \rangle : X, \}$ where $A_j \in A$ and $Collection \langle A_j \rangle \in S$.

To specify the type of data sets of a schema, parametric collection types are required. The issue of parametric types is a major one by itself. It will be elaborated only to a limited extent here by showing that a specific collection type (i.e., for a specific element type) is obtained from a parametric type by the same pushout construction [10] that we used for schema integration.

Definition 9 fObjc collection types) An instantiated collection interface $Collection \langle A \rangle$ is defined by the following pushout diagram

$$\begin{array}{ccc} T & \longrightarrow & Collection \langle T \rangle \\ \downarrow i & & \downarrow H \\ A & \xrightarrow{\wedge} & Collection \langle A \rangle \end{array}$$

In the above definition a parametric interface $Collection \langle T \rangle$ is viewed as a morphism $T \longrightarrow Collection \langle T \rangle$. The substitution $T \longrightarrow A$ and instantiation $Collection \langle T \rangle \xrightarrow{\wedge} Collection \langle A \rangle$ are also morphisms. The morphism $A \longrightarrow Collection \langle A \rangle$ is obtained from $T \longrightarrow Collection \langle T \rangle$ and the substitution $T \longrightarrow A$.

The above construction generalizes the usual views of instantiated parametric types. If a pushout is based on schema morphisms (which apply to both structural properties and integrity constraints), not just on signature morphisms, then the notion of instantiation of parametric types becomes semantic in nature.

Constraints are sentences expressed in Horn clause logic. Terms include method invocations and thus appear in the object-oriented form. Atoms are invocations of Boolean methods. This limited logic is sufficiently expressive to capture most typical database constraints, such as key and inclusion dependencies.

Definition 10 fObjc constraints) For a given schema signature *Sig* with sorts S :

- A collection object of type $Collection \langle A \wedge \rangle$ is a term of type $Collection \langle A \wedge \rangle$.
- A variable X of type $A \wedge$ where $A \wedge \in S$ is a term of type $A \wedge$.
- If a is a term of type $A \setminus$, a_2, \dots, a_n are terms of respective types $A \setminus, A_2, \dots, A_n$, $C \ m(A_i, A_2, \dots, A_n)$ is a method of the interface A_k of *Sig*, then $a.m(a_2, \dots, a_n)$ is a term of type C .
- With the above definitions, an atom is of the form $a.m(a_2, \dots, a_n)$ where m 's result type is Boolean.
- A constraint is of the form $p \leftarrow p_1, p_2, \dots, p_n$ where p, p_1, p_2, \dots, p_n are atoms.

Permissible schema signature transformations of **Objc** are defined below as schema signature morphisms. Their core is a mapping of sorts, extended to signatures of methods and collection objects. This mapping is the identity on the initial schema's predefined sorts, so these sorts have the same interpretation in all **Objc** schemas. Note that $Collection \langle T \rangle$ is not intended to have methods with arguments of type $Collection \langle T \rangle$, so that $Collection \langle A \wedge \rangle$ would be a structural subtype of $Collection \langle A \setminus \rangle$ if $A \setminus$ is a subtype of A_i .

Definition 11 *fObjc schema signature morphisms*) A morphism of schema signatures $\langle f \rangle : \text{Sigi} \rightarrow \text{Sig-z}$ consists of the following:

- An injective mapping of sorts $\langle f \rangle : S_1 \rightarrow S_2$ such that
 - $\langle f \rangle(s) = s$ for $s \in S_0$ where S_0 are the sorts of the initial schema Scho
 - $\langle f \rangle(\text{Collection} \langle A_i \rangle) = \text{Collection} \langle \langle f \rangle(A_i) \rangle$
- $\langle f \rangle$ applies to interfaces as follows: $\langle f \rangle(C \text{ m}(C_1, C_2, \dots, C_n)) = \langle f \rangle(C) (\langle f \rangle(m))(\langle f \rangle(C_1), \langle f \rangle(C_2), \dots, \langle f \rangle(C_n))$
- $\langle f \rangle$ maps collection objects of Sigi into collection objects of Sig-z such that $\langle f \rangle(\text{Collection} \langle A_j \rangle : X_j) = \langle f \rangle(\text{Collection} \langle A_j \rangle) : \langle f \rangle(X_j)$.

The above definition is intended to allow a mapping of an abstract data type to its extension with possible renaming and possibly enforcing structural subtyping conditions.

The notion of schema signature morphism is extended below to a schema morphism. The extension maps the integrity constraints in accordance with the mapping of the sorts in a schema signature morphism.

Definition 12 *fObjc schema morphisms*) A morphism of schemas $\langle cf \rangle : \text{Schi} \rightarrow \text{Sdi2}$ is a schema signature morphism $\langle f \rangle : \text{Sigi} \rightarrow \text{Sig2}$ such that $\langle j \rangle$ extends to a function $E_1 \rightarrow E_2$ as follows:

- $\langle j \rangle(xs) = X0(s)$
 $(s_1, s_2, \dots, s_n) = \langle cf \rangle(a). \langle cf \rangle(m) (\langle cf \rangle(a_1), \dots, \langle cf \rangle(a_n))$, where m may be a Boolean method.
- $\langle j \rangle(p \langle -pi, P_2, \dots, P_n \rangle) = \langle f \rangle(P) \langle f \rangle(P_1), \langle f \rangle(P_2), \dots, \langle f \rangle(P_n)$

Given Definitions 10 and 11, the definition of the *Sen* functor is immediate and so is the verification of its functorial properties.

Definition 13 *fObjc Sen functor*) Define $\text{Sen} : \text{Sign} \rightarrow \text{Set}$ as follows:

- $\text{Sen}(\text{Sig})$ is a set of sentences of the form $\{p \langle -pi, P_2, \dots, P_n \rangle\}$ where p, pi, P_2, \dots, P_n are atoms according to Definition 10.
- Given a schema signature morphism $\langle j \rangle : \text{Sigi} \rightarrow \text{Sig2}$,
 $\text{Sen}(\langle cf \rangle) (\langle f \rangle \langle -pi, P_2, \dots, P_n \rangle) = \langle f \rangle(P) \langle f \rangle(P_1), \langle f \rangle(P_2), \dots, \langle f \rangle(P_n)$

Proposition 1 *fObjc Sen functor*) Then Sen is a functor $\text{Sign} \rightarrow \text{Set}$.

Proof Sen obviously preserves the identity schema signature morphism, i.e. $\text{Sen}(\text{Id}_{\text{Sig}}) = \text{Id}_{\text{Sen}(\text{Sig})}$ and the composition of schema signature morphisms, i.e., $\text{Sen}(\langle f \rangle \langle j \rangle) = \text{Sen}(\langle f \rangle) \text{Sen}(\langle j \rangle)$

A database for a given **Objc** schema includes a domain for each sort in the schema. A domain is a set equipped with functions representing the interpretation of the method signatures. A database also includes the actual data sets, one for each collection object signature in the schema. By Definition 2, a database for a given schema must be consistent.

Definition 14 (*Objc databases*) A database d for a given schema $\text{Sch} = (\text{Sig}, E)$ consists of the following:

- A collection of sets (domains) $\{D_s \mid s \in S\}$ where S is the set of sorts of Sig . S includes Boolean, whose domain is true and false.
- For each method m of an interface A with the signature $C \text{ m}(C_1, C_2, \dots, C_n)$ a function $f_m : D_w \rightarrow D_c$ where $D_w = D_{A \text{ k}} \times D_{C_1} \times \dots \times D_{C_n}$ with $w = A \text{ k } C_1 \dots C_n$.
- For each collection variable of type $\text{Collection} \langle A \text{ k} \rangle$ a set with elements from the domain D_k .
- If $A_i \langle Z \text{ Au} \rangle$ then there exists a function (projection) $D_k \rightarrow D_i$ and thus also a function $\text{Collection} \langle A \text{ k} \rangle \rightarrow \text{Collection} \langle D_i \rangle$.
- Let e be a constraint with variables X such that $X_s \in D_s$ is a set of variables in X of sort s . If θ is a substitution of variables in e (i.e. a family of functions $X_s \rightarrow D_s$) then $e \theta$ evaluates to true.

Although collections in the above definition are interpreted as sets, in general they can be bags.

To complete the above definition, we must specify its morphisms of the category of databases: a family of functions, one per database domain, that map the actual data sets as well. These maps are required to satisfy two conditions: a standard algebraic condition that applies to operations on the original domains and their images, and a condition that applies to the integrity constraints. As we are talking here about the category of databases for a schema, not just for a schema signature, database morphisms are required to map each consistent database into another consistent database.

Definition 15 (Database morphisms for **Objc** schemas) Let d and d' be databases consistent with respect to a schema (Sig, E) . Then a database morphism $h : D \rightarrow D'$ is a family of functions $h_s : D_s \rightarrow D'_s$ for $s \in S$ where S stands for the set of sorts of Sig .

- For each method m of interface A with signature $C_m(C_1 x_1, C_2 x_2, \dots, C_n x_n)$ and a $G D_w$, $h_s(f_m(a)) = f_m(h_w(a))$ holds, where h_w is a product of functions $h_k : C_k x_k \rightarrow C_k h_c x_k$ when $w = A \wedge C_1 C_2 \dots C_n$ as in the following diagram:

$$\begin{array}{ccc} D_w & \xrightarrow{h_m} & D_c \\ h_w \downarrow & & \downarrow h_c \\ D'_w & \wedge & D'_c \end{array}$$

- Suppose $e \in E$ has variables X . If $\delta_s : X_s \rightarrow D_s$ is a substitution of variables for X and $e \in \delta$ evaluates to true, then so does $e \in \delta'$, where δ' is constructed as the composition of δ_s and h_s .

Definition 16 (Db functor for **Objc** schemas) Define Db as follows:

$Db(Sig)$ is a category with objects and arrows defined in 14 and 15. Given $\langle f \rangle : Sig_1 \rightarrow Sig_2$, $Db(\langle f \rangle)$ is defined as:

- If (f_e) is a Sig_2 database with domains $\{D_s \mid s \in S_2\}$ then $Db(\langle f \rangle)(d_2)$ is a Sig_1 database with domains $\{D_s \mid s \in S_1\}$. The same applies to the collection objects.
- A morphism $h : e_2 \rightarrow d'_2$ of Sig_2 databases (a family of functions $\{h_s \mid s \in S_2\}$), is mapped into a morphism $Db(\langle f \rangle)(h)$ of Sig_1 databases by restricting h to the family of functions $\{h_s \mid s \in S_1\}$.

Proposition 2 specifies how the Db functor is constructed and how its functorial properties are verified. Note that D_s and h_s now correspond to the sort $s \in S_1$.

Proposition 2 (Db functor for **Objc** schemas) Db in Definition 16 is a functor $Sign \rightarrow Cat^{op}$.

Proof A schema signature morphism $\langle f \rangle : Sig_1 \rightarrow Sig_2$ according to Definition 11 maps to a database morphism $Db(\langle f \rangle) : Db(Sig_2) \rightarrow Db(Sig_1)$ according to the above construction and Definition 15. The family $Db(\langle f \rangle)(d_2) \rightarrow Db(\langle f \rangle)(d'_2)$ is a Sig_1 database morphism, because h is a morphism of Sig_2 databases. Two properties are essential here: the injective mapping of sorts required in Definition 11 of schema signature morphisms and mapping of collections as specified in Definition 14.

The developments presented so far in this section lead to two theorems. The first one asserts that the paradigm satisfies the conditions of Definition 7 and thus is a schema transformation framework. It is based on the already established functors Sen and Db .

Theorem 2 (A schema transformation framework)

- Let the category of schema signatures **Objc** be defined according to Definitions 8 and 11.
- Define the functor $Sen : Sign \rightarrow Set$ according to Proposition 1.
- Define the functor $Db : Sign \rightarrow Cat^{op}$ according to the Proposition 2 so that the category $Db(Sig)$ is defined according to Definitions 14 and 15.
- The satisfaction relation is also specified in Definitions 14 and 15.

Under the above conditions the Integrity Requirement holds.

Proof We have to prove that each schema signature morphism $\langle f \rangle : Sig_A \rightarrow Sig_B$ in Definition 11 satisfies the Integrity Requirement, that is, $ds \models_{Sig_B} Sen(\langle j \rangle)(e)$ iff $Db(\langle f \rangle)(dB) \models_{Sig_A} e$ holds for each sentence $e \in Sen(Sig_A)$.

Given $e \in Sen(Sig_A)$, $Sen(\langle f \rangle)(e)$ is defined by Proposition 1. For a given object ds of the category $Db(Sig_B)$, $Db(\langle j \rangle)(d, B)$ is defined in Proposition 2. The proof amounts to:

$$dB \models_{Sig_B} \langle f \rangle(P) \iff Db(\langle f \rangle)(dB) \models_{Sig_A} (P \leftarrow Pi, Pi, \dots, Pn) \iff Db(\langle f \rangle)(dB) \models_{Sig_A} (P \leftarrow Pi, Pi, \dots, Pn) \blacksquare$$

The second theorem establishes the existence of the schema-join of any two schemas of this schema transformation framework. This result provides a standard technique for integrating two **Objc** schemas. This integration is both structural and semantic (i.e. the integrity constraints are properly integrated) because it is based entirely on schema morphisms according to the pushout construction in Definition 6.

Theorem 3 The schema-join of any two **Objc** schemas over the initial schema $Scho$ exists.

Proof Given schemas $Schi$ and $Schz$ the integrated schema $Schi * Sch^{\wedge}$ is defined as follows: (i) $Ssch! * Sch2 = \langle SiU5_2 \text{ (sorts)} \rangle$, (ii) $A_{Sch1} * Sch2 = A_{Sch1} \cup A_{Sch2}$ (interfaces), (in) $C_{Sch1} * Sch2 = C_{Sch1} \cup C_{Sch2}$ (collections), and (iv) $Esch! * Sch2 = Esch1 \cup Esch2$ (constraints).

With this construction $Schi$ and Sch^{\wedge} are subschemas of $Schi * Sch^{\wedge}$ since we have two injective schema morphisms $\langle f \rangle_i : Schi \rightarrow Schi * Sch^{\wedge}$ and $\langle f \rangle_2 : Sch,2 \rightarrow Schi * Sch,2$.

Suppose that we are given $p : Schi \rightarrow S12$ and $q : Sch,2 \rightarrow S2$. The required schema morphism $\langle j \rangle : Schi * Sch,2 \rightarrow Sch2$ is defined on the sorts as follows: $\langle f \rangle(s) = p(s)$ for $s \in Si$ and $\langle j \rangle(s) = q(s)$ for $s \in S2$.

Note that $\langle f \rangle_i(s) = s$ for $s \in So$ and $\langle fo \rangle(s) = s$ for $s \in So$. So^{\wedge} is well defined for $s \in Si \cap S2$.

As interfaces, collections and constraints of $Schi * Sch,2$ are defined as disjoint unions of the corresponding components of $Schi$ and $Sch,2$, we have $\langle f \rangle_h = p$ and $\langle j \rangle_k = q$.

Note that with the choice of a different logic the above result would not necessarily hold. The union of constraints could lead to a set of sentences for which there is no database (particularly in a given category) that satisfies the constraints (represents a model for the constraints) $\langle j \rangle_i \{Ei\}$ and $\wedge^{(\wedge)}$ (see Definition 3).

7 Conclusions and Related Work

The main contribution of this paper is a model theory for generic schema management. While the ideas on generic schema management proposed so far have been mostly informal [4] and pragmatic [5], this paper shows that a formal framework for such generic tools does indeed exist.

A distinctive feature of the presented paradigm is that it applies to a variety of categories of schemas and to a variety of logic bases for expressing database integrity constraints. This level of generality is accomplished by making use of a categorical model theory called *institutions*, proposed for programming language paradigms [9]. We believe ours is the first application of this theory to database models and languages. An earlier attempt in [15] is similar in spirit, but is less general. It also does not address the logic basis and preservation of integrity constraints and lacks provable results.

The core of this model theory is a general, transformation-oriented definition of a data model. This generic schema management framework serves as a pattern for constructing data models in such a way that schema transformations as well as the associated database transformations have well-defined meaning. As a rule, most well known data models have not been constructed in such a spirit. Furthermore, this framework is intended to be applied to data models that are still not completely or formally established, such as XML. A further distinctive feature of this framework is that it has a strong database integrity requirement as its possibly most fundamental component.

This paper also shows how to address some well-known problems such as schema equivalence [13, 14] and schema integration [6] at this level of generality. The results are independent of a particular category of schemas and apply to both structural properties and integrity constraints. The requirement for proper handling of the integrity constraints in those problems is one of the contributions of this paper. Furthermore, this approach is independent of a particular logic paradigm used as a basis for the constraint language.

Contrary to the published work on schema and data transformations (schema integration in particular) in which transformations at the meta and data levels have the same direction [19], this paper reveals a subtlety not considered in the relevant papers. The subtlety is manifested in the reversal of the transformation arrows at the two levels. This observation is an important component of the presented model theory, having both pragmatic and mathematical significance.

In this paper one option is that schema transformations must satisfy particular structural subtyping conditions. Such a strict discipline has its place in programming languages, but it is often not satisfied in schema transformations. Furthermore, if the abstract data types are equipped with constraints, semantic compatibility notions such as behavioral subtyping [16] become a major issue. Model theoretic implications of behavioral compatibility issues in mapping abstract data types equipped with constraints are given in

[1].

In a model in which inheritance is identified with subtyping (as in Java) a single partial order of sorts is required. This introduces further subtleties in schema transformations and schema integration. Models with different types of ordering of sorts are elaborated in [1] and [3]. The relevant results on pushouts of algebraic specifications are given in [12].

In this paper we do not consider the problem of mapping one schema transformation framework (for example, relational) into another (for example, object-oriented). This situation is captured by the notion of a morphism of schema transformation frameworks (following [9, 15]), and is a topic of work in progress.

The formal paradigm presented in this paper applies to appropriately simplified XML schemas [23]. The reasons are: the nature of the type system of the XML Schema, the kind of the integrity constraints (keys and referential integrity constraints) expressible in XML schemas (also considered in [8]), and the existence of the initial schema (more precisely, a name space). Space limitations do not allow further elaboration of these important implications of the development presented in this paper. However, application of this model theory to the XML data model is a major topic of future research.

This paper is not addressing explicitly modeling of database dynamics. But this is by no means a limitation of the paradigm. Database dynamics is taken into account by choosing the *Sen* functor for a suitable temporal logic. A development of such a paradigm is given in [1, 2, 3].

References

- [I] Alagic, S.: Semantics of temporal classes, *Information and Computation*, 163, pp. 60 - 102, 2000.
- [2] Alagic, S.: Constrained matching is type safe, Proceedings of the Sixth Int. Workshop on Database Programming Languages, *LNCS 1369*, Springer-Verlag, pp. 78-96, 1998.
- [3] Alagic, S. and M. Alagic, Order-sorted model theory for temporal executable specifications, *Theoretical Computer Science* 179, pp. 273-299, 1997.
- [4] Bernstein, P. A. , A. Halevy, R. A. Pottinger: A vision for management of complex models, *A CM SIGMOD Record* 29(4), 2000.
- [5] Bernstein, P.A. and E. Rahm: Data warehouse scenarios for model management. ER 2000, *LNCS 1920*, Springer-Verlag, pp. 1-15, 2000.
- [6] Batini, C, M. Lenzerini, and S. B. Navathe: A comparative analysis of methodologies for database schema integration, *ACM Computing Surveys* 18(4), pp. 323-364, 1986.
- [7] Buneman, P., S. Davidson, and A. Kosky: Theoretical aspects of schema merging. EDBT 1992, pp. 152-167.
- [8] Fan, W. and L. Libkin, On XML integrity constraints in the presence of DTDs, *Proc. of ACM PODS Conf.*, 2001.
- [9] Goguen, J. and R. Burstall, Institutions: Abstract model theory for specification and programming, *Journal of the ACM*, 39, No. 1, pp. 92-146, 1992.
- [10] Goguen, J.: Types as theories, in: G. M. Reed, A. W. Roscoe and R. F. Wachter, *Topology and Category Theory in Computer Science*, pp. 357-390, Clarendon Press, Oxford, 1991.
- [II] Goguen, J. and J. Meseguer, Order-sorted algebra I: Equational deduction for multiple inheritance, overloading, exceptions and partial operations, *Theoretical Computer Science* 105, pp. 217-273, 1992.
- [12] Haxthausen, A.E. and F. Nickl: Pushouts of order-sorted algebraic specifications, in M. Wirsing and M. Nivat, eds., *Algebraic Methodology and Software Tech. (AMAST 96)*, *LNCS 1101*, Springer-Verlag, pp. 132-147, 1996. [13] Hull, R.: Relative information capacity of simple relational database schemata, *SIAM Journal of Computing*, Vol. 15, No. 3, pp.856 -886, 1986. [14] Hull, R.: Managing semantic heterogeneity in databases: A theoretical perspective, *PODS 1997*, 51-61 [15] Kalinichenko, L.A.: Methods and tools for equivalent data model mapping construction, *Proc. of EDBT*, *LNCS 416*, Springer-Verlag, pp. 92-119, 1990. [16] Liskov, B. and J. M. Wing: A behavioral notion of subtyping, *ACM TOPLAS* 16, pp. 1811 - 1841. [17] Mac Lane, S.: *Categories for a Working Mathematician*, Springer, 1998.
- [18] Madhavan, J., P. A. Bernstein, E. Rahm: Generic schema matching with Cupid, *VLDB 01*, to appear. [19] Miller, R. J., Y. E. Ioannidis, R. Ramakrishnan, Schema equivalence in heterogeneous systems: Bridging theory and practice, *Information Systems Vol. 19*, No. 1, pp. 3-31, 1994. [20] Rahm, E. and P.A. Bernstein: On matching schemas automatically. MSR Tech. Report MSR-TR-2001-17, 2001. [21] Spaccapietra, S., C. Parent, Y. Dupont: Model independent assertions for integration of heterogeneous schemas, *VLDB Journal* 1, pp. 81-126, 1992. [22] Spaccapietra, S. and C. Parent: View integration: A step forward in solving structural conflicts, *IEEE TKDE* 6(2), pp. 258 - 274, 1994. [23] W3C: XML Schema, <http://www.w3c.org/XML/schema>, 2001.