

Algebraic Model Management: A Survey

Patrick Schultz, David Spivak, Ryan Wisnesky

► **To cite this version:**

Patrick Schultz, David Spivak, Ryan Wisnesky. Algebraic Model Management: A Survey. 23th International Workshop on Algebraic Development Techniques (WADT), Sep 2016, Gregynog, United Kingdom. pp.56-69, 10.1007/978-3-319-72044-9_5. hal-01767471

HAL Id: hal-01767471

<https://hal.inria.fr/hal-01767471>

Submitted on 16 Apr 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Algebraic Model Management: A Survey

Patrick Schultz¹, David I. Spivak¹, and Ryan Wisnesky²

¹ Massachusetts Institute of Technology, Cambridge, MA USA

² Categorical Informatics, Inc., Cambridge, MA USA
ryan@catinf.com

Abstract. We survey the field of model management and describe a new model management approach based on algebraic specification.

Keywords: Model Management · Algebraic Databases · Data Integration · Functorial Data Migration · Category Theory · Algebraic Specification

1 Introduction

In this paper we survey the field of *model management* and describe a new model management approach based on techniques from the field of *algebraic specification*, with the hope of establishing an interlingua between the two fields. By “model management” we mean “meta-data intensive” database management in the sense of Bernstein & Melnik [4], which we define in Section 2. By “a new algebraic model management approach” we mean our particular way [19] [20] of specifying database schemas and instances using algebraic (equational) theories.

We first noticed a connection between model management and algebraic specification while investigating applications of *category theory* [3] to *data integration* [5]. These investigations are described in [19] and [20], and we present no substantial new results in this paper. We assume readers have basic proficiency with category theory [3], algebraic specification [17], and SQL.

Outline. In Section 2 we describe the traditional approach to model management and in Section 3 we describe our algebraic approach. Also in Section 3 we describe the open-source AQL (Algebraic Query Language) tool, available for download at <http://categoricaldata.net/aql.html>, which implements our approach in software. We conclude in Section 4 by comparing our approach with the traditional approach.

2 Model Management

To quote from Melnik [16]:

Many challenging problems facing information systems engineering involve the manipulation of complex metadata artifacts, or *models*, such as database schemas, interface specifications, or object diagrams, and

mappings between models. The applications that solve metadata manipulation problems are complex and hard to build. The goal of generic model management is to reduce the amount of programming needed to develop such applications by providing a database infrastructure in which a set of high-level algebraic operators, such as Match, Merge, and Compose, are applied to models and mappings as a whole rather than to their individual building blocks.

In the paragraph above the word “model” is defined to mean a metadata artifact such as a schema, which conflicts with the definition of the word “model” as a structure satisfying a theory. In this paper, we use the phrase “model management” to mean the field identified above, and use the word “model” to mean a structure satisfying a theory.

Today model management is a large sub-field of information management with a research literature containing hundreds of published articles [4]. There is a consensus in that literature [4] that model management is concerned with at least the problems described in the next sections.

2.1 Schema mapping

Given two database schemas S and T , the *schema mapping* problem [7] is to construct a “mapping” $F : S \rightarrow T$ that captures some user-specified relationship between S and T . Different model management systems use different notions of schema, including SQL, XML, and RDF [4]. The most common mapping formalism studied in the literature is that of “embedded dependencies” (EDs) [5]: formulae in a fragment of first-order logic with useful computational properties.

We will use SQL schemas and EDs in our examples in this section. Consider the following SQL schema S , consisting of two tables connected by a foreign key:

```
CREATE TABLE N2(ID INT PRIMARY KEY, age INT)
```

```
CREATE TABLE N1(ID INT PRIMARY KEY, name STRING, salary INT,
  f INT FOREIGN KEY REFERENCES N2(ID))
```

and the following SQL schema T , consisting of one table:


```
CREATE TABLE N(ID INT PRIMARY KEY, age INT,
  name STRING, salary INT).
```

These two SQL schemas are displayed graphically in Figure 1.

An example schema mapping $F : S \rightarrow T$ expressing that the target table N is the join of source tables $N1$ and $N2$ along the column f is:

$$\forall id_1, id_2, a, n, s. N1(id_1, n, s, id_2) \wedge N2(id_2, a) \rightarrow N(id_1, a, n, s).$$

Two instances satisfying the above ED are shown in Figure 1. In general, many EDs can map between two SQL schemas.



2.2 Query generation

Given a schema mapping $F : S \rightarrow T$, the *query generation* problem [4] is to construct a query which converts databases on S to databases on T in a way that satisfies F . The query languages typically studied include SQL, XQuery, and various comprehension- and λ -calculi [4].

A SQL query to implement the example mapping from Section 2.1 is:

```

INSERT INTO N
SELECT N1.ID, N1.age, N2.name, N1.sal
FROM N1, N2
WHERE N1.f = N2.ID

```

Technically, the INSERT portion of the above SQL code is not a “query”, but rather an “update”, and in practice the code generated from a query generation task will often store the results of the query. An example of running the above SQL is shown as the left-to-right direction of Figure 1. In general, many or no SQL queries may implement a set of EDs [5]. EDs can also be directly executed by an algorithm called “the chase” [5].

2.3 Mapping Inversion

Given a schema mapping $F : S \rightarrow T$, the *mapping inversion* problem [10] is to construct a schema mapping $F^{-1} : T \rightarrow S$ that undoes F with respect to query generation (i.e. the queries generated from F and F^{-1} should be inverses).

The natural candidate ED to invert the schema mapping of Section 2.1 expresses that N projects onto N1 and N2:

$$\forall id_1, a, n, s. N(id_1, a, n, s) \rightarrow \exists id_2. N1(id, n, s, id_2) \wedge N2(id_2, a)$$

and a possible SQL implementation of this ED is:

<pre>INSERT INTO N1 SELECT ID, name, sal, ID FROM N</pre>	<pre>INSERT INTO N2 SELECT ID, age FROM N</pre>
---	---

However, the above ED is *not* an inverse to the ED of Section 2.1, as is seen by taking $\emptyset = N1 \neq N2$. Indeed, it is rare for an ED, or set of EDs, to be invertible, and weaker notions of inverse, such as “quasi-inverse” [10], are common in the literature [10]. An example of running the above SQL is shown as the right-to-left direction of Figure 1.

2.4 Mapping Composition

Given schema mappings $F : S \rightarrow T$ and $G : T \rightarrow U$, the *mapping composition* problem [8] is to construct a schema mapping $G \circ F : S \rightarrow U$ that is equivalent with respect to the query generation problem (i.e. running the query generated from $G \circ F$ should have the same effect as running the query generated from G on the results of the query generated from F).

The composition of the ED from Section 2.1 with the ED from Section 2.3 is

$$\forall id_1, id_2, n, s, a. N1(id_1, n, s, id_2) \wedge N2(id_2, a) \rightarrow \exists x. N1'(id, n, s, x) \wedge N2'(x, a)$$

where $N1'$, $N2'$ are target “copies” of source tables $N1$, $N2$. This composed ED is not the identity, thereby showing that the ED from Section 2.3 does not invert the ED from Section 2.1. In the case of EDs, composed mappings may not exist [8], but some restrictions and extensions of EDs are closed under composition [8].

2.5 Schema matching

Given two database schemas S and T , the *schema matching* problem [5] is to automatically find “correspondences” between S and T and to automatically infer schema mappings $S \rightarrow T$ from these correspondences. In general, inference of entire mappings cannot be fully automated and the focus of the matching problem is to reduce the human effort required to construct a schema mapping by e.g., suggesting partial mappings that can be completed by users. There are many techniques for schema matching ranging from comparison of column names by string similarity to machine learning algorithms; for an overview, see [5]. In the example from Section 2.1, two correspondences that are easy to automatically find are $(N1, N)$ and $(N2, N)$ and tools such as Clio [14] can create the ED from Section 2.1 from these two correspondences.

2.6 Further References

In this paper we will focus on the problems described in the previous sections, but many other problems are studied in the model management literature [4], and many of these problems are related to algebraic specification. For example, *schema/instance merge* problems [4], which arise often in data integration scenarios [2], can be formalized as pushouts in suitable categories of schemas/instances [20], and such pushouts are related to model-theoretic concepts such as *model amalgamation* [15].

Many software products solve model management problems [4], including *ETL (Extract, Transform, Load)* tools [5], which extract data from separate databases, apply user-specified transformations, and then load the result into a target system such as a data warehouse; *query mediators* [5], which answer queries about a “virtual” integrated database by combining queries about separate source databases; and *visual schema mapping tools* [14] which allow users to create schema mappings by visually connecting related schema elements with lines, as shown in Figure 2.

There have been at least two attempts to provide a “meta semantics” for model management operations. In [16] Melnik gives a “state based” meta semantics to some of the above operations by defining a schema mapping $S \rightarrow T$ to be an arbitrary binary relation between instances on S and instances on T ; the ED-based semantics described above is an instantiation of this meta semantics. In [2] and [13] the authors give an “institution theoretic” meta semantics to some of the above operations by defining a schema mapping $S \rightarrow T$ to be a morphism in a suitable category of schemas; AQL’s semantics is an instantiation of this meta semantics.

3 Algebraic Model Management

Our approach to model management is based on the algebraic approach to databases, data migration, and data integration we describe in [19] and [20]. Those works, and hence this work, extend a particular category-theoretic data model that originated in the late 1990s [11] and was later extended in [21] and [23] and implemented in AQL (<http://categoricaldata.net/aql.html>).

In the next section we describe our formalism for database schemas and instances and introduce AQL. The subsequent sections implement the model management operations from Section 2 using our formalism. In this section we abbreviate “algebraic theory” as “theory”.

3.1 Algebraic Databases

In our formalism [20], database schemas and instances are defined as theories of a certain kind, which we describe in the next sections. For ease of exposition, we will sometimes conflate schemas and instances as defined in our formalism with their AQL equivalents.




Fig. 2. A schema mapping in Clío [14]

Type sides We first fix a theory, Ty , called the *type side* of our formalism. The sorts of Ty are called *types* and the functions of Ty are the functions that can appear in schemas and instances.

AQL allows arbitrary theories to be used as type sides. But we have found that in practice, AQL users almost always want to use the theory of an existing programming language, say java, for their type side. The ability to “bind” AQL to an existing language is particularly important in model management because input data may only be accessible through, e.g., a java API. For this reason, AQL allows a type side to be defined by specifying, for each sort s , a java class C_s and a java function $\text{String} \rightarrow C_s$ that tells AQL how to interpret the strings it encounters in AQL programs as objects of C_s .

An example AQL type side about integers and strings is shown in Figure 3. This type side defines a theory with two sorts and infinitely many constants – all the java strings and integers – and no equations. The java code for `Int` says that whenever a string x is encountered in an AQL program and a term of sort `Int` is required, that java’s `parseInt` function should be applied to x to yield the desired `Int`. The keyword `literal`, used in many places in AQL, indicates a literal (user-defined constant) definition.

Schemas A *schema* on type side Ty is a theory extending Ty with new sorts (called *entities*), new unary functions from entities to types (called *attributes*), new unary functions from entities to entities (called *foreign keys*), and new equa-

```

typeside Ty = literal {
  java_types
    String = "java.lang.String"
    Int = "java.lang.Integer"
  java_constants
    String = "return input[0]"
    Int = "return java.lang.Integer.parseInt(input[0])"
}

```

Fig. 3. AQL type side Ty

<pre> schema S = literal : Ty { entities N1 N2 foreign_keys f : N1 -> N2 attributes name : N1 -> String salary : N1 -> Int age : N2 -> Int } </pre>	<pre> schema T = literal : Ty { entities N attributes name : N -> String salary : N -> Int age : N -> Int } </pre>
---	---

Fig. 4. AQL schemas S and T on type side Ty

tions (called *data integrity constraints*) of the form $\forall v : s. t = t'$, where s is an entity and t, t' are terms of the same type, each containing a single free variable v . The restrictions in the preceding sentence (e.g., no functions from types to entities) are necessary to use our formalism for model management purposes [19] [20].

Figure 4 shows the AQL schemas corresponding to Figure 1. These schemas contain no equations and are both on the type side Ty defined in Figure 3.

Instances An *instance* I on schema S is a theory extending S with new 0-ary function (constant) symbols called *generators* and non-quantified equations. An example AQL instance on schema S (Figure 4) is shown in Figure 5.

The intended meaning of an instance I , written $\llbracket I \rrbracket$, is the *term model* (i.e., *initial algebra*) for I which contains, for each sort s , a *carrier set* consisting of the closed terms of sort s modulo provability in I . A morphism of instances $I \rightarrow J$ is a homomorphism (natural transformation) of algebras $\llbracket I \rrbracket \rightarrow \llbracket J \rrbracket$.

Figure 6 shows the meaning of the instance I from Figure 5 in the AQL tool. The AQL tool visually displays term models as sets of tables, one per entity e , each with an ID column corresponding to the carrier set for e . The tables in Figure 6 are isomorphic to the left tables in Figure 1.


```

instance I = literal : S {
  generators
    1 2 3 : N1
  equations
    name(1) = Alice   salary(1) = 100   age(f(1)) = 20
    name(2) = Bob     salary(2) = 250   age(f(2)) = 20
    name(3) = Sue     salary(3) = 300   age(f(3)) = 30
}

```

Fig. 5. AQL instance I on schema S

N1 (3)				N2 (3)	
ID	name	salary	f	ID	age
[1]	Alice	100	[1.f]	[1.f]	20
[2]	Bob	250	[2.f]	[2.f]	20
[3]	Sue	300	[3.f]	[3.f]	30

Fig. 6. Initial algebra for AQL instance I

```

mapping F = literal : S -> T {
  entities
    N1 -> N
    N2 -> N
  foreign_keys
    f -> lambda x:N. x
  attributes
    name -> lambda x:N. name(x)
    salary -> lambda x:N. salary(x)
    age -> lambda x:N. age(x)
}

```

Fig. 7. AQL schema mapping $F : S \rightarrow T$

In the following sections we implement the model management operations from Section 2 using the preceding definitions of schema and instance.

3.2 Schema mapping

Given schemas S, T , the *schema mapping* problem (Section 2.1) is to construct a “mapping” $F : S \rightarrow T$ that captures some relationship between S and T .

Let S and T be AQL schemas on the same type side Ty . An AQL schema mapping $F : S \rightarrow T$ is defined as a “derived signature morphism” [18] from S to T that is the identity on Ty . That is, $F : S \rightarrow T$ assigns to each entity $e \in S$ an entity $F(e) \in T$, and to each attribute / foreign key $f : s \rightarrow s'$ a term $F(f)$, of type $F(s')$ and with one free variable of type $F(s)$, in a way that respects equality: if $S \vdash t = t'$, then $T \vdash F(t) = F(t')$. We have found that many mappings arising in practice cannot be expressed using plain signature morphisms and require the more general notion of “derived” signature morphism.

Whereas a schema mapping in Section 2.1 was an ED (formula in a fragment of first-order logic), which induces a single binary satisfaction relation between instances, AQL schema mappings are derived signature morphisms and induce three relations between instances, which we will describe in the next section.

An example AQL schema mapping $F : S \rightarrow T$ is shown in Figure 7, where AQL schemas S and T are defined in Figure 4. This mapping is also shown graphically in Figure 1.

3.3 Query generation

Given a mapping $F : S \rightarrow T$, the *query generation* problem (Section 2.2) is to use F to construct a query which converts databases on S to databases on T .

In our formalism, the database instances and morphisms on a schema S constitute a category, denoted $S\text{-Inst}$, and a schema mapping $F : S \rightarrow T$ induces a functor $\Sigma_F : S\text{-Inst} \rightarrow T\text{-Inst}$ defined by substitution. The functor Σ_F has a right adjoint, $\Delta_F : T\text{-Inst} \rightarrow S\text{-Inst}$, which corresponds to the “model reduct functor” when our formalism is described in institution-theoretic terms [2]. The functor Δ_F has a right adjoint, $\Pi_F : S\text{-Inst} \rightarrow T\text{-Inst}$. See [19] for proof that Δ_F always has left and right adjoints. As adjoints, Δ_F, Π_F preserve limits and Δ_F, Σ_F preserve colimits, implying many useful properties; for example, $\Sigma_F(I + J) \cong \Sigma_F(I) + \Sigma_F(J)$ and $\Pi_F(I \times J) \cong \Pi_F(I) \times \Pi_F(J)$.

Note that unlike Section 2.1, where there was a single query associated with a schema mapping (ED), in our algebraic approach there are three queries, one for each of $\Delta_F, \Sigma_F, \Pi_F$. The conditions under which $\Delta_F, \Sigma_F, \Pi_F$ can be expressed in SQL and vice-versa are characterized in [23].

Although it is possible to give explicit formulae to define $\Delta_F, \Sigma_F, \Pi_F$ [19] we instead give examples in Figures 1 and 8. Note that in these examples we are not showing instances (theories) as defined in Section 3.1; we are showing term models. For this reason, we surround $\Delta_F, \Sigma_F, \Pi_F$ with denotation brackets $\llbracket \ \rrbracket$ in these examples. In addition, as adjoints Δ, Σ, Π are only defined up to unique

isomorphism, so we arbitrarily make up names for IDs and in these examples. Figures 1 and 8 show an AQL schema mapping F which takes two distinct source entities, N1 and N2, to the target entity N. The $[\Delta_F]$ functor projects in the opposite direction of F : it projects columns from the single table for N to two separate tables for N1 and N2, similar to `FROM N AS N1` and `FROM N AS N2` in SQL. When there is a foreign key from N1 to N2, the $[\Delta_F]$ functor populates it so that N can be recovered by joining N1 and N2. The $[\Pi_F]$ functor takes the cartesian product of N1 and N2 when there is no foreign key between N1 and N2, and joins N1 and N2 along the foreign key when there is. The $[\Sigma_F]$ functor disjointly unions N1 and N2; because N1 and N2 are not union compatible (have different columns), $[\Sigma_F]$ creates null values. When there is a foreign key between N1 and N2, $[\Sigma_F]$ merges the tuples that are related by the foreign key, resulting in a join. As these examples illustrate, Δ_F can be thought of as projection, Π_F can be thought of as a product followed by a filter (which can result in a join), and Σ_F can be thought of as a disjoint union (which does not require union-compatibility) followed by a merge (which can also result in a join).

3.4 Mapping Composition


Given schema mappings $F : S \rightarrow T$ and $G : T \rightarrow U$, the *mapping composition* problem (Section 2.4) is to construct a schema mapping $G \circ F : S \rightarrow U$ that is equivalent with respect to query generation.

In one sense, the mapping composition problem is trivial [19] for our formalism: $\Delta_{F \circ G} \cong \Delta_G \circ \Delta_F$, $\Pi_{F \circ G} \cong \Pi_F \circ \Pi_G$, and $\Sigma_{F \circ G} \cong \Sigma_F \circ \Sigma_G$. But this solution is not wholly satisfactory because in practice a mixture of Δ , Σ , Π functors may be needed to accomplish any particular task (similarly, in SQL a mixture of joins and unions may be needed to accomplish any particular task). The following results are proved in [19] and [23]:

- Every composition $\Sigma_F \circ \Delta_G$ is isomorphic to $\Delta_{F'} \circ \Sigma_{G'}$ for some F', G' . This statement is also true if Σ_F is replaced with Π_F .
- Pairs of the form (F, G) , denoting $\Sigma_F \circ \Delta_G$, are closed under composition. This statement is also true if Σ_F is replaced with Π_F . Such pairs can be specified in an intuitive “select-from-where” syntax, described in [19] and [20].
- Triples of the form (F, G, H) , denoting $\Sigma_F \circ \Pi_G \circ \Delta_H$, are closed under composition, provided that F is a *discrete op-fibration* [3], which is exactly the “union compatibility” condition [5] that Σ_F performs unions over tables whose columns match; Figure 1 is not a discrete op-fibration.

3.5 Mapping Inversion

Given a schema mapping $F : S \rightarrow T$, the *mapping inversion* problem (Section 2.3) is to construct a mapping $F^{-1} : T \rightarrow S$ that somehow “undoes” F .



<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: center; width: 33%;">N1</td> <td style="text-align: center; width: 33%;">N2</td> <td style="width: 34%;"></td> </tr> <tr> <td style="border-bottom: 1px solid black;">ID name salary</td> <td style="border-bottom: 1px solid black;">ID age</td> <td style="text-align: center;">$\leftarrow [\Delta_F]$</td> </tr> <tr> <td>1 Alice \$100</td> <td>4 20</td> <td style="border-bottom: 1px solid black;">ID name salary age</td> </tr> <tr> <td>2 Bob \$250</td> <td>5 20</td> <td>1 Alice \$100 20</td> </tr> <tr> <td>3 Sue \$300</td> <td>6 30</td> <td>2 Bob \$250 20</td> </tr> <tr> <td></td> <td></td> <td>3 Sue \$300 30</td> </tr> </table>	N1	N2		ID name salary	ID age	$\leftarrow [\Delta_F]$	1 Alice \$100	4 20	ID name salary age	2 Bob \$250	5 20	1 Alice \$100 20	3 Sue \$300	6 30	2 Bob \$250 20			3 Sue \$300 30																				
N1	N2																																					
ID name salary	ID age	$\leftarrow [\Delta_F]$																																				
1 Alice \$100	4 20	ID name salary age																																				
2 Bob \$250	5 20	1 Alice \$100 20																																				
3 Sue \$300	6 30	2 Bob \$250 20																																				
		3 Sue \$300 30																																				
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: center; width: 33%;">N1</td> <td style="text-align: center; width: 33%;">N2</td> <td style="width: 34%;"></td> </tr> <tr> <td style="border-bottom: 1px solid black;">ID name salary</td> <td style="border-bottom: 1px solid black;">ID age</td> <td style="text-align: center;">$\xrightarrow{[\Sigma_F]}$</td> </tr> <tr> <td>1 Alice \$100</td> <td>4 20</td> <td style="border-bottom: 1px solid black;">ID name salary age</td> </tr> <tr> <td>2 Bob \$250</td> <td>5 20</td> <td>1 Alice \$100 age(1)</td> </tr> <tr> <td>3 Sue \$300</td> <td>6 30</td> <td>2 Bob \$250 age(2)</td> </tr> <tr> <td></td> <td></td> <td>3 Sue \$300 age(3)</td> </tr> <tr> <td></td> <td></td> <td>4 name(4) salary(4) 20</td> </tr> <tr> <td></td> <td></td> <td>5 name(5) salary(5) 20</td> </tr> <tr> <td></td> <td></td> <td>6 name(6) salary(6) 30</td> </tr> </table>	N1	N2		ID name salary	ID age	$\xrightarrow{[\Sigma_F]}$	1 Alice \$100	4 20	ID name salary age	2 Bob \$250	5 20	1 Alice \$100 age(1)	3 Sue \$300	6 30	2 Bob \$250 age(2)			3 Sue \$300 age(3)			4 name(4) salary(4) 20			5 name(5) salary(5) 20			6 name(6) salary(6) 30											
N1	N2																																					
ID name salary	ID age	$\xrightarrow{[\Sigma_F]}$																																				
1 Alice \$100	4 20	ID name salary age																																				
2 Bob \$250	5 20	1 Alice \$100 age(1)																																				
3 Sue \$300	6 30	2 Bob \$250 age(2)																																				
		3 Sue \$300 age(3)																																				
		4 name(4) salary(4) 20																																				
		5 name(5) salary(5) 20																																				
		6 name(6) salary(6) 30																																				
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: center; width: 33%;">N1</td> <td style="text-align: center; width: 33%;">N2</td> <td style="width: 34%;"></td> </tr> <tr> <td style="border-bottom: 1px solid black;">ID name salary</td> <td style="border-bottom: 1px solid black;">ID age</td> <td style="text-align: center;">$\xrightarrow{[\mathcal{I}_F]}$</td> </tr> <tr> <td>1 Alice \$100</td> <td>4 20</td> <td style="border-bottom: 1px solid black;">ID name salary age</td> </tr> <tr> <td>2 Bob \$250</td> <td>5 20</td> <td>1 Alice \$100 20</td> </tr> <tr> <td>3 Sue \$300</td> <td>6 30</td> <td>2 Bob \$250 20</td> </tr> <tr> <td></td> <td></td> <td>3 Sue \$300 20</td> </tr> <tr> <td></td> <td></td> <td>4 Alice \$100 20</td> </tr> <tr> <td></td> <td></td> <td>5 Bob \$250 20</td> </tr> <tr> <td></td> <td></td> <td>6 Sue \$300 20</td> </tr> <tr> <td></td> <td></td> <td>7 Alice \$100 30</td> </tr> <tr> <td></td> <td></td> <td>8 Bob \$250 30</td> </tr> <tr> <td></td> <td></td> <td>9 Sue \$300 30</td> </tr> </table>	N1	N2		ID name salary	ID age	$\xrightarrow{[\mathcal{I}_F]}$	1 Alice \$100	4 20	ID name salary age	2 Bob \$250	5 20	1 Alice \$100 20	3 Sue \$300	6 30	2 Bob \$250 20			3 Sue \$300 20			4 Alice \$100 20			5 Bob \$250 20			6 Sue \$300 20			7 Alice \$100 30			8 Bob \$250 30			9 Sue \$300 30		
N1	N2																																					
ID name salary	ID age	$\xrightarrow{[\mathcal{I}_F]}$																																				
1 Alice \$100	4 20	ID name salary age																																				
2 Bob \$250	5 20	1 Alice \$100 20																																				
3 Sue \$300	6 30	2 Bob \$250 20																																				
		3 Sue \$300 20																																				
		4 Alice \$100 20																																				
		5 Bob \$250 20																																				
		6 Sue \$300 20																																				
		7 Alice \$100 30																																				
		8 Bob \$250 30																																				
		9 Sue \$300 30																																				

Fig. 8. Example Data Migrations (see Section 3.2)

Our formalism has strong inversion properties but does not have inverses per se. When there exists $F^{-1} : T \rightarrow S$ such that $F \circ F^{-1} = id$ and $F^{-1} \circ F = id$, then $\Delta_F \circ \Delta_{F^{-1}} \cong id$, $\Sigma_F \circ \Sigma_{F^{-1}} \cong id$, and $\Pi_F \circ \Pi_{F^{-1}} \cong id$. In general F need not have an inverse, but when S and T have finite initial algebras / term models (which is a priori undecidable, and implies decidability of S and T) it is possible to construct F^{-1} whenever it exists by considering all possible functors $T \rightarrow S$. When F has a right adjoint $G : T \rightarrow S$, a weaker condition than having an inverse, there are canonical morphisms $\Sigma_F \rightarrow \Delta_G$ and $\Delta_F \rightarrow \Pi_G$.

In practice “round-tripping” [5] of data is desirable even when inverses do not exist. For example, projection, because it forgets information, typically cannot be inverted, but we may want to remember where the projected data originated. In our formalism the adjunctions between Σ, Δ, Π provide round-tripping. For example, for every $F : S \rightarrow T$ and S -instance I there is a canonical morphism $I \rightarrow \Delta_F(\Sigma_F(I))$, the *unit* of the $\Sigma_F \dashv \Delta_F$ adjunction, which describes where each ID in I is sent to by Σ_F (and similarly for Π_F). Dually, for every T -instance J there is a canonical morphism $\Sigma_F(\Delta_F(J)) \rightarrow J$, the *co-unit* of the $\Sigma_F \dashv \Delta_F$ adjunction, which describes where the IDs in $\Delta_F(J)$ originate (and similarly for Π_F). The unit and co-unit can be used to obtain, for every morphism $h : \Sigma_F(I) \rightarrow J$, a mate $h' : I \rightarrow \Delta_F(J)$ and vice-versa (and similarly for Π_F). Relating adjointness to existing relaxed notions of inverse such as quasi-inverse [9] is an important area for future work.

3.6 Schema matching

Given database schemas S and T , the *schema matching* problem (Section 2.5) is to automatically suggest schema mappings $S \rightarrow T$ to the user.

In this section, we define two schema matching techniques used by AQL. Our techniques compare entities, and foreign keys and attributes (“symbols”) by name, as strings, and so our techniques depend on having (probably user-provided) names whose similarity as strings reflects their semantic similarity. Let $\sigma : \text{String}, \text{String} \rightarrow [0, 1]$ be any string similarity function [5] where a value of 1 indicates a “good” match and a value of 0 indicates a “bad” match.

- The first technique attempts to infer a schema mapping $F : S \rightarrow T$. For each entity $s \in S$, we define $F(s) := t$ where $t \in T$ is an entity that maximizes $\sigma(s, t)$. For each symbol $f : s \rightarrow s' \in S$, we then consider the set X of symbols $F(s) \rightarrow F(s')$. If X is non-empty, we choose a symbol $g \in X$ that maximizes $\sigma(f, g)$ and set $F(f) := g$. If X is empty but there is a shortest path p from $F(s)$ to $F(s')$, we set $F(f) := p$. If no shortest path exists, the match fails. The F so constructed is only a candidate schema mapping: AQL must verify that F preserves provable equality in S .
- The second technique attempts to infer a schema A and schema mappings $F : A \rightarrow S$ and $G : A \rightarrow T$. Such a span of mappings can be interpreted as a query of the form $\Sigma_F \circ \Delta_G$ or $\Pi_F \circ \Delta_G$. Let c be some user-provided string similarity cutoff. The entities of A are those pairs of S -entities and

T -entities (s, t) such that $\sigma(s, t) > c$. The symbols $(s, t) \rightarrow (s', t')$ of A are those pairs of S -symbols and T -symbols $(f : s \rightarrow s', g : t \rightarrow t')$ such that $\sigma(f, g) > c$. The mappings F and G are projections.

4 Conclusion

When comparing our algebraic approach to model management with other approaches originating in relational database theory [1] it is important to note that our databases are “deductive databases” [1]. That is, we define databases “intentionally”, as sets of equations, rather than as sets of tables. As such, care must be taken when mediating between our definitions and relational definitions. For example, our instances can be “inconsistent” in the sense that an instance can prove $1 = 2$ for two distinct constant symbols 1 and 2. Such situations are often, but not always [12], errors, and the AQL tool checks for such situations using standard techniques based on “conservative theory extensions” [12]. In addition, our schemas do not define a set of constants (a “domain”) that all the instances on that schema share, as is customary in relational database theory [7]. Hence our approach is closer in spirit to traditional logic [6] than database theory [1].

There are many connections between our algebraic approach to model management and the ED-based approach described in Section 2. EDs are more expressive than our purely equational data integrity constraints and can be added to our formalism in a simple way, described in [22] (although in [22], EDs are called “lifting problems”). In ED-based approaches the “chase” [5] operation has a similar semantics to our Σ operation, and a formal comparison between the chase and Σ is forthcoming. **Acknowledgements.** The authors thank Lucian Popa, Eswaran Subrahmanian, and Peter Gates and were supported by NIST SBIR grant 70NANB 16H178, AFOSR grant FA9550-14-1-0031 and NASA grant NNL14AA05C.

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley-Longman (1995)
2. Alagic, S., Bernstein, P.: A model theory for generic schema management. DBPL (2001)
3. Barr, M., Wells, C.: Category Theory for Computing Science. Prentice Hall International (1995)
4. Bernstein, P.A., Melnik, S.: Model management 2.0: Manipulating richer mappings. ICMD (2007)
5. Doan, H., Halevy, A., Ives, Z.: Principles of Data Integration. Morgan Kaufmann (2012)
6. Enderton, H.B.: A Mathematical introduction to logic. Academic Press (2001)
7. Fagin, R., Kolaitis, P.G., Miller, R.J., Popa, L.: Data exchange: Semantics and query answering. Theoretical Computer Science (2005)

8. Fagin, R., Kolaitis, P.G., Popa, L., Tan, W.: Composing schema mappings: Second-order dependencies to the rescue. *TODS* (2005)
9. Fagin, R., Kolaitis, P.G., Popa, L., Tan, W.: Quasi-inverses of schema mappings. *TODS* (2008)
10. Fagin, R.: Inverting schema mappings. *TODS* (2007)
11. Fleming, M., Gunther, R., Rosebrugh, R.: A database of categories. *Journal of Symbolic Computation* 35(2) (2003)
12. Ghilardi, S., Lutz, C., Wolter, F.: Did I damage my ontology? *Principles of Knowledge Representation and Reasoning* (2006)
13. Goguen, J.: Information integration in institutions (unpublished). <http://cseweb.ucsd.edu/~goguen/pps/ifi04.pdf> (2004)
14. Haas, L.M., Hernández, M.A., Ho, H., Popa, L., Roth, M.: Clio grows up: From research prototype to industrial tool. *ICMD* (2005)
15. Hodges, W.: *A Shorter Model Theory*. Cambridge University Press (1997)
16. Melnik, S.: *Generic Model Management: Concepts And Algorithms* (Lecture Notes in Computer Science). Springer-Verlag (2004)
17. Mitchell, J.C.: *Foundations of Programming Languages*. MIT Press (1996)
18. Mossakowski, T., Krumnack, U., Maibaum, T.: What is a derived signature morphism? *RTADT* (2014)
19. Schultz, P., Spivak, D.I., Vasilakopoulou, C., Wisnesky, R.: *Algebraic databases. Theory and Applications of Categories* (2017)
20. Schultz, P., Wisnesky, R.: Algebraic data integration (unpublished). <http://arxiv.org/abs/1503.03571> (2016)
21. Spivak, D.I.: Functorial data migration. *Information and Computation* (2012)
22. Spivak, D.I.: Database queries and constraints via lifting problems. *Mathematical Structures in Computer Science* (2014)
23. Spivak, D.I., Wisnesky, R.: Relational foundations for functorial data migration. *DBPL* (2015)