

# *Algebraic data integration\**

PATRICK SCHULTZ

*Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA*  
(e-mail: schultzp@mit.edu)

RYAN WISNESKY

*Categorical Informatics, Inc., Cambridge, MA, USA*  
(e-mail: ryan@catinf.com)

---

## Abstract

In this paper, we develop an algebraic approach to data integration by combining techniques from functional programming, category theory, and database theory. In our formalism, database schemas and instances are algebraic (multi-sorted equational) theories of a certain form. Schemas denote categories, and instances denote their initial (term) algebras. The instances on a schema  $S$  form a category,  $S\text{-Inst}$ , and a morphism of schemas  $F : S \rightarrow T$  induces three adjoint data migration functors:  $\Sigma_F : S\text{-Inst} \rightarrow T\text{-Inst}$ , defined by substitution along  $F$ , which has a right adjoint  $\Delta_F : T\text{-Inst} \rightarrow S\text{-Inst}$ , which in turn has a right adjoint  $\Pi_F : S\text{-Inst} \rightarrow T\text{-Inst}$ . We present a query language based on for/where/return syntax where each query denotes a sequence of data migration functors; a pushout-based design pattern for performing data integration using our formalism; and describe the implementation of our formalism in a tool we call AQL (Algebraic Query Language).

---

## 1 Introduction

In this paper, we develop an algebraic approach to data integration by combining techniques from functional programming, category theory, and database theory. By *data integration* we mean combining separate but related database schemas and instances to form a coherent, unified schema and instance, and we consider query and data migration to be special cases of data integration. By *algebraic* we mean that our schemas are algebraic (purely equational) theories and our instances denote algebras (models) of our schemas. We use category theory to define the semantics of our approach: schemas and instances form categories, and data integration operations are characterized with categorical constructions such as adjoint functors and pushouts. We use techniques from functional programming to implement our approach by constructing syntactic objects and reasoning about them, both on paper and using automated techniques. We use database theory as a baseline with which to compare our approach.

\* The authors would like to thank David Spivak and Peter Gates. Patrick Schultz was supported by AFOSR grant FA9550-14-1-0031, ONR grant N000141310260, and NASA grant NNH13ZEA001N. Ryan Wisnesky was supported by NIST SBIR grant 70NANB15H290.

The mathematics of the semantics of our approach are worked out in detail in Schultz *et al.* (2017) using sophisticated category theory. In this paper, our goal is to implement a syntax for this semantics, and crucially, to do so in a *computable* way. How to do this is not obvious, because the mathematical objects defined in Schultz *et al.* (2017) are almost always infinite and not computable. This paper is a comprehensive description of the implementation of the AQL tool (<http://categoricaldata.net/aql.html>), an open-source data integration tool capable of solving problems similar to those solved by relational data integration tools such as Clio (Haas *et al.*, 2005) and Rondo (Melnik *et al.*, 2003), as well as query languages such as SQL and LINQ (Grust, 2004).

Because our approach draws on functional programming, category theory, and database theory, the more knowledge a reader has about each of these fields, the more the reader will get out of the paper. These three theories are used in a deep, rather than wide, way: we use mostly basic categorical concepts such as category, functor, natural transformation, and adjunction; we use mostly basic functional programming concepts such as equational logic and algebraic data types; and we use mostly basic database theory concepts such as conjunctive queries and labeled nulls. For this reason, we believe that a reader well-versed in only one of these areas can still get something out of this paper, and will be rewarded with a deeper insight into the other areas, at least from a data integration perspective. We include short primers on category theory (Section 2) and equational logic (Section 3), and connections to database theory are made as remarks in the text.

### 1.1 Background

Our data model extends a particular category-theoretic data model that we call the *functorial data model* (FDM) (Spivak, 2012). Originating in the late 1990s (Fleming *et al.*, 2003), the FDM defines a schema  $S$  to be a finite presentation of a category (Barr & Wells, 1995): a directed, labeled multi-graph and a set of path equality constraints as shown in Figure 1. In this figure, the nodes of the graph indicate “entities” of employees, departments, and strings, and the arrows represent functions between entities, such as the assignment of a name to each employee. The equations in the graph serve as data integrity constraints, indicating, for example, that secretaries work in the same department they are the secretary for. An instance  $I$  on  $S$  (a.k.a., an  $S$ -instance) is a set of tables, one per node in  $S$ , and a set of columns, one per edge in  $S$ , that obey the equations in  $S$ . Such an instance determines a functor  $S \rightarrow \mathbf{Set}$ , where  $\mathbf{Set}$  is the category whose objects are sets and whose morphisms are functions.

In the FDM, the database instances on a schema  $S$  (i.e., functors  $S \rightarrow \mathbf{Set}$ ) constitute a category, denoted  $S\text{-Inst}$ , and a functor (a.k.a. schema mapping; Fagin *et al.* 2005b)  $F : S \rightarrow T$  between schemas  $S$  and  $T$  induces three adjoint data migration functors:  $\Delta_F : T\text{-Inst} \rightarrow S\text{-Inst}$ , defined as  $\Delta_F(I) := I \circ F$  (note that  $F : S \rightarrow T$  and  $I : T \rightarrow \mathbf{Set}$ ), and the left and right adjoints to  $\Delta_F$ , respectively:  $\Sigma_F : S\text{-Inst} \rightarrow T\text{-Inst}$  and  $\Pi_F : S\text{-Inst} \rightarrow T\text{-Inst}$ . The  $\Sigma, \Delta, \Pi$  data migration functors provide a category-theoretic alternative to the traditional relational operations for

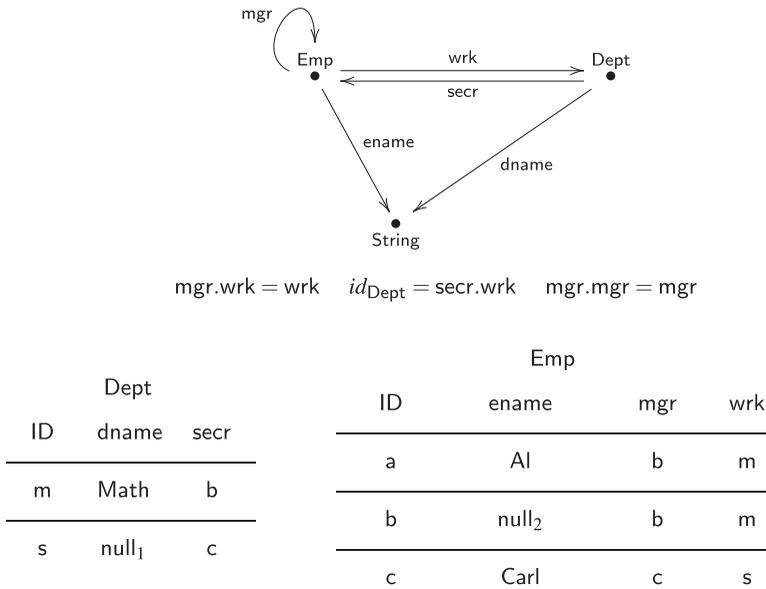
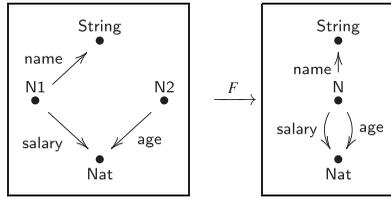


Fig. 1. A schema and instance in the functorial data model.

both querying data (relational algebra) and migrating/integrating data (“chasing” embedded dependencies (EDs); Fagin *et al.*, 2005b). Their relative advantages and disadvantages over the relational operations are still being studied, but see Spivak (2012) for a preliminary discussion. At a high-level,  $\Delta$  can be thought of as a projection,  $\Pi$  as join, and  $\Sigma$  as union.

An example schema mapping  $F : S \rightarrow T$  is shown in Figure 2. A full description of this figure is given in Section 4.2; here we sketch an overview of the figure to give intuition about  $\Delta_F, \Pi_F$ , and  $\Sigma_F$ . The functor  $F : S \rightarrow T$  is the unique edge-preserving map from  $S$  to  $T$ . The  $\Delta_F$  operation takes the table  $N$  and projects it to two tables. The  $\Pi_F$  operation performs a cartesian product of  $N_1$  and  $N_2$ , and the  $\Sigma_F$  operation performs an “outer join” (Garcia-Molina *et al.*, 2008) of  $N_1$  and  $N_2$ ; i.e., it unions two tables that have different columns by adding null values as necessary. When there is an edge  $f : N_1 \rightarrow N_2$  (Figure 9),  $\Delta_F$  performs a “lossless join decomposition” (Garcia-Molina *et al.*, 2008) along  $f$  of  $N$ ,  $\Pi_F$  performs a join along  $f$ , and  $\Sigma_F$  performs a union-then-merge (Roberson & Wyss, 2004), resulting in a join for the particular instance in the figure.

The FDM’s basic idea of schemas-as-categories and three adjoint data migration functors  $\Delta, \Sigma, \Pi$  recurs in our data model, but we base our formalism entirely on algebraic (equational) logic and therefore diverge from the original FDM. We define database schemas and instances to be equational theories of a certain kind. A schema mapping  $F : C \rightarrow D$  is defined as a morphism (provability-respecting translation) of equational theories  $C$  and  $D$ , and we define the  $\Sigma_F$  data migration as substitution along  $F$ . The conditions we impose on our equational theories guarantee that  $\Sigma_F$  has a right adjoint,  $\Delta_F$ , which in turn has a right adjoint,  $\Pi_F$ . In practice, programming with  $\Delta, \Sigma$ , and  $\Pi$  is verbose, and so we define a terse query language where each



N1			N2		N			
ID	Name	Salary	ID	Age	ID	Name	Salary	Age
1	Alice	\$100	1	20	1	Alice	\$100	20
2	Bob	\$250	2	20	2	Bob	\$250	20
3	Sue	\$300	3	30	3	Sue	\$300	30

$\xleftarrow{[\Delta_F]}$

N1			N2		N			
ID	Name	Salary	ID	Age	ID	Name	Salary	Age
1	Alice	\$100	1	20	a	Alice	\$100	1.age
2	Bob	\$250	2	20	b	Bob	\$250	2.age
3	Sue	\$300	3	30	c	Sue	\$300	3.age
					d	4.name	4.salary	20
					e	5.name	5.salary	20
					f	6.name	6.salary	30

$\xrightarrow{[\Sigma_F]}$

N1			N2		N			
ID	Name	Salary	ID	Age	ID	Name	Salary	Age
1	Alice	\$100	1	20	a	Alice	\$100	20
2	Bob	\$250	2	20	b	Bob	\$250	20
3	Sue	\$300	3	30	c	Sue	\$300	20
					d	Alice	\$100	20
					e	Bob	\$250	20
					f	Sue	\$300	20
					h	Alice	\$100	30
					i	Bob	\$250	30
					j	Sue	\$300	30

$\xrightarrow{[\Pi_F]}$

Fig. 2. Example functorial data migrations.

query is a collection of generalized for/where/return expressions (Abiteboul *et al.*, 1995). Each query can be evaluated into a data migration of the form  $\Delta \circ \Pi$  (and vice-versa) and each query can be “co-evaluated” into a data migration of the form

$\Delta \circ \Sigma$  (and vice-versa). To integrate data we must go beyond  $\Delta, \Sigma, \Pi$ , and we use pushouts (Barr & Wells, 1995) of schemas and instances as the basis for a data integration design pattern suitable for building data warehouses.

## 1.2 Outline

The rest of this paper is divided into five sections. In Section 2, we review category theory, and in Section 3 we review algebraic (multi-sorted equational) logic. In Section 4, we describe how we use algebraic theories to define schemas, instances, and the other artifacts necessary to perform data migration and query. In Section 5, we describe how our formalism is implemented in the AQL tool. In Section 6, we describe how we use pushouts of schemas and instances to perform data integration and include an extended example.

## 1.3 Related work

In this section, we describe how our work relates to the functional data model (Section 1.3.1), functional programming (Section 1.3.2), relational data integration (Section 1.3.3), and other work that treats schemas as categories (Section 1.3.4). Mathematical related work is discussed in Schultz *et al.* (2017).

### 1.3.1 Versus the functional data model

Our data model is simultaneously an extension of, and a restriction of, the functional data model (Shipman, 1981). Both formalisms use functions, rather than relations, over entities and data types as their principal data structure. The two primary ways the functional data model extends ours is that it allows products of entities (e.g., a function  $\text{Person} \times \text{Department} \rightarrow \text{String}$ ), and it allows some non-equational data integrity constraints (e.g., a constraint  $f(x) = f(y) \rightarrow x = y$ ). These two features can be added to our data model as described in Spivak (2014), although doing so alters our data model's properties (e.g., the existence of  $\Pi$  cannot in general be guaranteed in the presence of non-equational constraints). Therefore, our data model can be thought of as extending the fragment of the functional data model where schemas are categories. Restricting to schemas-as-categories allows us to extend the functional data model with additional operations (e.g.,  $\Pi$ ) and to provide strong static reasoning principles (e.g., eliminating the need for run time data integrity checking; see Section 4.3.1). The functional data model also includes updates, a topic we are still studying. The schemas of our data model and the functional data model both induce graphs, so both data models are graph data models in the sense of Angles and Gutierrez (2008). See Spivak (2012) for a discussion of how our formalism relates to RDF.

### 1.3.2 Versus functional programming

Our formalism and functional programming languages both extend equational logic (Mitchell, 1996). As discussed in Section 4.1, many artifacts of our formalism,

including schemas and user-defined functions, are “algebraic data types” and our instances denote implementations (algebras) of these types. (Our use of “algebraic data type” must be understood in the algebraic specification sense, as a data type specified by a set of equations (Mitchell, 1996), rather than in the Haskell/ML sense of a data type specified by products and sums.) As such, the implementation of our formalism draws heavily on techniques from functional programming: for example, when a set of data integrity constraints forms a confluent and terminating rewrite system, our implementation uses reduction to normal forms to decide equality under the constraints (Section 5.1). Our formalism’s use of types, rather than terms, to represent sets (e.g., entity sets such as `Person`) is common in expressive type theories such as Coq (Bertot & Castéran, 2010) and contrasts with the sets-as-terms approach used in simply-typed comprehension calculi (Grust, 2004); for a comparison of the two approaches (which we have dubbed QINL and LINQ, respectively), see Schultz *et al.* (2015). Certain categories of schemas and instances are cartesian closed (Barr & Wells, 1995), meaning that certain schemas and instances can be defined by expressions in the simply-typed  $\lambda$ -calculus, and meaning that certain categories of schemas and instances are models of the simply-typed  $\lambda$ -calculus, but we have yet to find an application of this fact.

### 1.3.3 Versus relational data integration

Our formalism is an alternative to the traditional relational formalisms for both querying data (relational algebra) and migrating/integrating data (“chasing” EDs; Fagin *et al.*, 2005b). In Spivak & Wisnesky (2015), we proved that the  $\Delta, \Sigma, \Pi$  operations can express any (select, project, cartesian product, union) relational algebra query and gave conditions for the converse to hold. Our formalism uses purely equational data integrity constraints, which can be captured using first-order EDs, but not all first-order EDs can be captured using only equations. However, our formalism can be extended in a simple way, described in Spivak (2014), to capture all first-order EDs; when this is done, we find that a parallel chase step of an ED on an instance can be expressed using the pushout construction of Section 6. A proof that our formalism, extended as in Spivak (2014) and with a fixed point operation, is at least as expressive as chasing with first-order EDs is forthcoming. At present, we do not understand the relationship between our formalism and second-order EDs (Fagin *et al.*, 2005a), although our “uber-flower” queries (Section 4.3) can be written as second-order EDs of a particular form (Section 4.3.6). For additional discussion about how our formalism relates to EDs and the field of “model management” (Melnik *et al.*, 2003), see the survey paper by Schultz *et al.* (2016).

The duality between  $\Sigma$  and  $\Pi$  in our formalism suggests a “missing” operator in relational algebra: the dual to join. Because a join is a product followed by a filter along an equivalence relation (an equalizer), its dual is a sum (disjoint union) followed by a tuple-merge along an equivalence relation (a co-equalizer). The tuple-merge operation appears in relational data integration settings (Roberson & Wyss, 2004), and so does the idea of forming equivalence classes of tuples, in the guise

of the chase (Fagin *et al.*, 2005a), but we have not seen the dual to join explicitly described as such in the relational literature.

Our formalism defines databases “intensionally,” as sets of equations, and so in relational parlance our databases are “deductive databases” (Abiteboul *et al.*, 1995). As such, some care must be taken when mediating between relational definitions and categorical definitions. For example, our instances can be inconsistent, in the sense that an instance might prove that  $1 = 2$  for two distinct constant symbols 1 and 2, and only consistent instances can be meaningfully translated into relational instances. 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 data integration (Fagin *et al.*, 2005b). For these and other reasons (mentioned throughout this paper), our work is closer in spirit to traditional logic (Enderton, 2001) than database theory (Doan *et al.*, 2012).

The pushout data integration pattern (Section 6) is a “global as view” (Doan *et al.*, 2012) pattern, because the integrated schema is a function (the pushout) of the source schemas. But rather than relating the integrated schema to the source schemas by EDs or queries, we use functors. A pushed-out instance satisfies a universal property similar to that of a universal solution to a set of EDs (Fagin *et al.*, 2005b). Pushouts are investigated for relational data integration purposes in Alagic & Bernstein (2001). In that paper, the authors describe a design pattern for data integration that applies to a large class of formalisms: the so-called *institutions* (Goguen & Burstall, 1984). Our formalism is an institution, but their work differs from ours in key respects. First, they are primarily concerned with the  $\Delta$  data migration functor (they call our  $\Delta$  functor “Db” in their paper), because  $\Delta$  exists in all institutions. They recognize that pushouts (what they call “schema joins”) are a canonical way of obtaining integrated schemas, and that not all institutions have pushouts of schemas (ours does). Their main theorem is that in any institution the  $\Delta$  functor can be used to migrate the data on a pushout schema back to the source schemas. Our design pattern uses the  $\Sigma$  functor to go the other way: pushing data from source schemas to the integrated schema. See Goguen (2004) for more information about data integration in institutions. In the more general setting of algebraic specification, pushouts have received considerable attention as a means to integrate specifications (Blum *et al.*, 1987).

#### 1.3.4 Versus the functorial data model

Our work is related to a family of data models that treat database schemas as categories or variations thereof (Johnson *et al.*, 2002; Fleming *et al.*, 2003; Spivak, 2012, 2014; Spivak & Wisnesky, 2015; Schultz *et al.*, 2017). We refer to these data models as “functorial data models.” The original functorial data model (Fleming *et al.*, 2003) treated schemas as finitely presented categories and instances as set-valued functors, but is difficult to use for data integration purposes because most constructions on set-valued functors are only defined up to unique isomorphism, and in the context of data integration, we must distinguish two kinds of values in a database: atomic values such as strings and integers that must be preserved by morphisms (e.g., Bill), and meaningless identifiers that need not be preserved (e.g.,

ID	Name	Age	Salary		ID	Name	Age	Salary		ID	Name	Age	Salary
1	Alice	20	\$100	$\cong$ (good)	4	Alice	20	\$100	$\cong$ (bad)	1	Amy	20	\$100
2	Bob	20	\$250		5	Bob	20	\$250		2	Bill	20	\$250
3	Sue	30	\$300		6	Sue	30	\$300		3	Susan	30	\$300

Fig. 3. The attribute problem.

auto-generated IDs). (In contexts outside of data integration, such as query, there may not be a need to distinguish two types of values.) For example, the situation in Figure 3, which holds in the original functorial data model, is untenable for data integration purposes.

Several approaches to this “attribute problem” have been proposed, including Johnson *et al.* (2002) and Spivak & Wisnesky (2015). This paper extends the latter paper by defining database schemas to exist in an ambient computational context called a “type-side.” Data values that inhabit types (e.g., Bill : String) are preserved by database morphisms, but other values, such as meaningless identifiers, are not. As a result, our formalism does not suffer from the attribute problem, and solving the attribute problem was a significant motivation for our work.

Recently, Patterson (2017) defined “relational ologs,” and in so doing generalized the functorial data model (where instances are functors  $C \rightarrow \mathbf{Set}$ ) to relations: Patterson’s instances are functors  $C \rightarrow \mathbf{Rel}$ , where  $\mathbf{Rel}$  is the category of sets and relations. Interestingly, over finite databases, Patterson’s formalism, which is graphical in nature, is equivalent to the extension of our formalism by EDs (see Section 1.3.3). The AQL tool’s built-in example “FOAF” (Friend of a Friend) describes this equivalence. Independently, Zielinski *et al.* (2013) proposed the “Allegorical Conceptual Data Model,” a similar relational generalization of the functorial data model.

## 2 Review of category theory

In this section, we review standard material on category theory (Barr & Wells, 1995). Readers familiar with category theory can safely skim or skip this section.

### 2.1 Categories with products

A category  $\mathbf{C}$  consists of a class of objects  $Ob(\mathbf{C})$  and a class of morphisms  $Hom(\mathbf{C})$  between objects. Each morphism  $m$  has a source object  $S$  and a target object  $T$ , which we write as  $m : S \rightarrow T$ . Every object  $X$  has an identity morphism  $id_X : X \rightarrow X$ . When  $X$  is clear from the context we will write  $id_X$  as simply  $id$ . Two morphisms  $f : B \rightarrow C$  and  $g : A \rightarrow B$  may be composed, written  $f \circ g : A \rightarrow C$  or  $g ; f : A \rightarrow C$ . Composition is associative and  $id$  is its unit:

$$f \circ id = f \quad id \circ f = f \quad f \circ (g \circ h) = (f \circ g) \circ h$$

A morphism  $f : X \rightarrow Y$  is an *isomorphism* when there exists a  $g : Y \rightarrow X$  such that

$$f \circ g = id \quad g \circ f = id$$

Two objects are *isomorphic* when there exists an isomorphism between them. Example categories include:

- **Set**, the *category of sets*. The objects of **Set** are sets, and a morphism  $f : X \rightarrow Y$  is a (total) function from set  $X$  to set  $Y$ . Given morphisms  $f : Y \rightarrow Z$  and  $g : X \rightarrow Y$ , the morphism  $f \circ g : X \rightarrow Z$  is defined as function composition:  $(f \circ g)(x) := f(g(x))$ . The isomorphisms of **Set** are bijective functions. For each object  $X$ ,  $id_X$  is the identity function on  $X$ .
- Any directed graph generates a category, called the *free* category on the graph: its objects are the vertices of the graph, and its morphisms are the paths in the graph. For each vertex  $X$ ,  $id_X$  is the 0-length path  $X \rightarrow X$ . Composition of morphisms is concatenation of paths, and there are no non-identity isomorphisms.

A category **C** is said to *have products* when for every pair of objects  $X, Y$  in **C**, there exists an object  $X \times Y$  in **C**, morphisms  $\pi_1 : X \times Y \rightarrow X$  and  $\pi_2 : X \times Y \rightarrow Y$  in **C**, and for every pair of morphisms  $f : A \rightarrow X, g : A \rightarrow Y$  in **C**, there exists a morphism  $\langle f, g \rangle : A \rightarrow X \times Y$  in **C**, such that

$$\pi_1 \circ \langle f, g \rangle = f \quad \pi_2 \circ \langle f, g \rangle = g$$

and such that for every morphism  $h : A \rightarrow X \times Y$ ,

$$\langle \pi_1 \circ h, \pi_2 \circ h \rangle = h$$

For example, set-theoretic cartesian product is a product in the category of sets, **Set**.

## 2.2 Functors

A *functor*  $F : \mathbf{C} \rightarrow \mathbf{D}$  between two categories **C** and **D** is a mapping of objects of **C** to objects of **D** and morphisms of **C** to morphisms of **D** that preserves identities and composition:

$$F(f : X \rightarrow Y) : F(X) \rightarrow F(Y) \quad F(id_X) = id_{F(X)} \quad F(f \circ g) = F(f) \circ F(g)$$

Example functors include the following:

- For any category **C**, the identity functor  $Id : \mathbf{C} \rightarrow \mathbf{C}$  maps each object and morphism to itself.
- For any categories **C** and **D** and  $D$  an object of **D**, there exists a constant functor taking each object  $C$  in **C** to  $D$  and each morphism in **C** to  $id_D$ .
- The power set functor  $\mathcal{P} : \mathbf{Set} \rightarrow \mathbf{Set}$  maps each set to its power set and each function  $f : X \rightarrow Y$  to the function that sends  $U \subseteq X$  to its image  $f(U) \subseteq Y$ .
- For each set  $A$ , there is a product functor  $- \times A : \mathbf{Set} \rightarrow \mathbf{Set}$  mapping each set  $X$  to the cartesian product  $X \times A$  and each function  $f : X \rightarrow Y$  to the function  $f \times id_A : X \times A \rightarrow Y \times A$ .

A functor  $F : \mathbf{C} \rightarrow \mathbf{Set}$  for any category **C** is said to be *set-valued*.

### 2.3 Natural transformations

A natural transformation  $\alpha : F \Rightarrow G$  between two functors  $F : \mathbf{C} \rightarrow \mathbf{D}$  and  $G : \mathbf{C} \rightarrow \mathbf{D}$  is a family of morphisms  $\alpha_X : F(X) \rightarrow G(X)$  in  $\mathbf{D}$ , one for each object  $X$  in  $\mathbf{C}$ , such that for every  $f : X \rightarrow Y$  in  $\mathbf{C}$ ,

$$\alpha_Y \circ F(f) = G(f) \circ \alpha_X$$

This equation may conveniently be expressed as a *commutative diagram*:

$$\begin{array}{ccc} F(X) & \xrightarrow{F(f)} & F(Y) \\ \alpha_X \downarrow & & \downarrow \alpha_Y \\ G(X) & \xrightarrow{G(f)} & G(Y) \end{array}$$

A natural transformation  $\alpha$  is a *natural isomorphism* when for every object  $X$  in  $\mathbf{C}$ , the morphism  $\alpha_X$  is an isomorphism in  $\mathbf{D}$ . Example natural transformations include the following:

- The identity natural isomorphism  $id_F : F \Rightarrow F$  for a functor  $F : \mathbf{C} \rightarrow \mathbf{D}$  is defined as  $id_{F_X} : F(X) \rightarrow F(X) := id_{F(X)}$ .
- Consider the power set functor  $\mathcal{P} : \mathbf{Set} \rightarrow \mathbf{Set}$ . There is a natural transformation  $sng : Id_{\mathbf{Set}} \Rightarrow \mathcal{P}$  that maps every set  $X$  to the singleton set  $\{X\}$  (i.e.,  $sng_X : X \rightarrow \mathcal{P}(X)$ ), and there is a natural transformation  $union : \mathcal{P} \circ \mathcal{P} \Rightarrow \mathcal{P}$  that maps a set of sets  $\{X_1, \dots, X_n\}$  to its  $n$ -ary union  $X_1 \cup \dots \cup X_n$  (i.e.,  $union_X : \mathcal{P}(\mathcal{P}(X)) \rightarrow \mathcal{P}(X)$ ).

### 2.4 Adjunctions

An *adjunction* between categories  $\mathbf{C}$  and  $\mathbf{D}$  consists of a functor  $F : \mathbf{D} \rightarrow \mathbf{C}$  called the left adjoint, a functor  $G : \mathbf{C} \rightarrow \mathbf{D}$  called the right adjoint, a natural transformation  $\epsilon : F \circ G \Rightarrow Id_{\mathbf{C}}$  called the counit, and a natural transformation  $\eta : Id_{\mathbf{D}} \Rightarrow G \circ F$  called the unit, such that for every object  $X$  in  $\mathbf{C}$  and  $Y$  in  $\mathbf{D}$ , the following equations hold:

$$id_{F(Y)} = \epsilon_{F(Y)} \circ F(\eta_Y) \quad id_{G(X)} = G(\epsilon_X) \circ \eta_{G(X)}$$

Consequently, the set of morphisms  $F(Y) \rightarrow X$  is bijective with the set of morphisms  $Y \rightarrow G(X)$ . Example adjunctions include the following:

- Let  $A$  be a set and consider the product functor  $- \times A : \mathbf{Set} \rightarrow \mathbf{Set}$ . The exponential functor  $-^A : \mathbf{Set} \rightarrow \mathbf{Set}$ , which maps each set  $X$  to the set of functions from  $A$  to  $X$  (written  $X^A$ ), is right adjoint to  $- \times A$ . Intuitively, this is because the set of functions  $X \times Y \rightarrow Z$  is bijective with the set of functions  $X \rightarrow Z^Y$ .
- Consider the category of groups and group homomorphisms,  $\mathbf{Grp}$ . The functor  $free : \mathbf{Set} \rightarrow \mathbf{Grp}$ , which maps each set  $X$  to the free group generated by  $X$ , and the functor  $forget : \mathbf{Grp} \rightarrow \mathbf{Set}$ , which maps each group to its underlying set, are adjoint. Intuitively, maps from the free group  $free(X)$  to a group  $Y$  correspond

precisely to maps from the set  $X$  to the set  $\text{forget}(Y)$ : each homomorphism from  $\text{free}(X)$  to  $Y$  is fully determined by its action on generators.

### 3 Review of multi-sorted equational logic

In this section, we review standard material on multi-sorted equational logic, following the exposition in Mitchell (1996). Theories in multi-sorted equational logic are also called “algebraic theories,” as well as “Lawvere theories” and “product theories.” We will use these phrases interchangeably. For a category-theoretic study of such theories, see Adamek *et al.* (2011). Readers familiar with equational logic can safely skim or skip this section. We will write “theory” instead of “multi-sorted equational theory” in this section.

#### 3.1 Syntax

In this section, we define the syntax of theories. Readers may wish to refer to the example theory about strings and natural numbers in Figure 4 while reading this section.

A signature  $\text{Sig}$  consists of

1. a set  $\text{Sorts}$  whose elements are called *sorts*;
2. a set  $\text{Symbols}$  of pairs  $(f, s_1 \times \dots \times s_k \rightarrow s)$  with  $s_1, \dots, s_k, s \in \text{Sorts}$  and no  $f$  occurring in two distinct pairs. We write  $f : X$  instead of  $(f, X) \in \text{Symbols}$ . When  $k = 0$ , we may call  $f$  a *constant symbol* and write  $f : s$  instead of  $f : \rightarrow s$ . Otherwise, we may call  $f$  a *function symbol*.

We assume we have some countably infinite set  $\{v_1, v_2, \dots\}$ , whose elements we call *variables* and which are assumed to be distinct from any sort or symbol we ever consider. A *context*  $\Gamma$  is defined as a finite set of variable-sort pairs, with no variable given more than one sort:

$$\Gamma := \{v_1 : s_1, \dots, v_k : s_k\}$$

When the sorts  $s_1, \dots, s_k$  can be inferred, we may write a context as  $\{v_1, \dots, v_k\}$ . We may write  $\{v_1 : s, \dots, v_k : s\}$  as  $\{v_1, \dots, v_k : s\}$ . We may write  $\Gamma \cup \{v : s\}$  as  $\Gamma, v : s$ . We inductively define the set  $\text{Terms}^s(\text{Sig}, \Gamma)$  of *terms* of sort  $s$  over signature  $\text{Sig}$  and context  $\Gamma$  as

1.  $x \in \text{Terms}^s(\text{Sig}, \Gamma)$ , if  $x : s \in \Gamma$ ;
2.  $f(t_1, \dots, t_k) \in \text{Terms}^s(\text{Sig}, \Gamma)$ , if  $f : s_1 \times \dots \times s_k \rightarrow s$  and  $t_i \in \text{Terms}^{s_i}(\text{Sig}, \Gamma)$  for  $i = 1, \dots, k$ . When  $k = 0$ , we may write  $f$  for  $f()$ . When  $k = 1$ , we may write  $t_1.f$  instead of  $f(t_1)$ . When  $k = 2$ , we may write  $t_1 f t_2$  instead of  $f(t_1, t_2)$ .

We refer to  $\text{Terms}^s(\text{Sig}, \emptyset)$  as the set of *ground* terms of sort  $s$ . We will write  $\text{Terms}(\text{Sig}, \Gamma)$  for the set of all terms in context  $\Gamma$ , i.e.,  $\bigcup_s \text{Terms}^s(\text{Sig}, \Gamma)$ . Substitution of a term  $t$  for a variable  $v$  in a term  $e$  is written as  $e[v \mapsto t]$  and is recursively defined as usual:

$$v[v \mapsto t] = t \quad v'[v \mapsto t] = v' \quad (v \neq v') \quad f(t_1, \dots, t_n)[v \mapsto t] = f(t_1[v \mapsto t], \dots, t_n[v \mapsto t])$$

$Sorts := \{\text{Nat}, \text{Char}, \text{String}\}$   
 $Symbols := \{\text{zero} : \text{Nat}, \text{succ} : \text{Nat} \rightarrow \text{Nat}, A : \text{Char}, B : \text{Char}, \dots, Z : \text{Char},$   
 $\text{nil} : \text{String}, \text{cons} : \text{Char} \times \text{String} \rightarrow \text{String}, \text{length} : \text{String} \rightarrow \text{Nat}, + : \text{String} \times \text{String} \rightarrow \text{String}\}$   
 $Equations := \{\text{length}(\text{nil}) = \text{zero}, \forall c : \text{Char}, s : \text{String}. \text{length}(\text{cons}(c, s)) = \text{succ}(\text{length}(s))\}$

We will abbreviate zero as 0 and  $\text{succ}^n(\text{zero})$  as n. We will write e.g. Bill for

$\text{cons}(B, \text{cons}(I, \text{cons}(L, \text{cons}(L, \text{nil}))))$ .

Fig. 4. The multi-sorted equational theory *Type*.

$$\begin{array}{c}
 \frac{t \in \text{Terms}^s(\text{Sig}, \Gamma)}{Th \vdash \forall \Gamma. t = t : s} \qquad \frac{Th \vdash \forall \Gamma. t = t' : s}{Th \vdash \forall \Gamma. t' = t : s} \qquad \frac{Th \vdash \forall \Gamma. t = t' : s \quad Th \vdash \forall \Gamma. t' = t'' : s}{Th \vdash \forall \Gamma. t = t'' : s} \\
 \\
 \frac{Th \vdash \forall \Gamma. t = t' : s \quad v \notin \Gamma}{Th \vdash \forall \Gamma, v : s'. t = t' : s} \qquad \frac{Th \vdash \forall \Gamma, v : s. t = t' : s' \quad Th \vdash \forall \Gamma. e = e' : s}{Th \vdash \forall \Gamma. t[v \mapsto e] = t'[v \mapsto e'] : s'}
 \end{array}$$

Fig. 5. Inference rules for multi-sorted equational logic.

We will only make use of substitutions that are sort-preserving; i.e., to consider  $e[v \mapsto t]$ , we require  $e \in \text{Terms}^s(\text{Sig}, \Gamma)$  for some  $\Gamma$  such that  $v : s \in \Gamma$  and  $t \in \text{Terms}^s(\text{Sig}, \Gamma)$ . To indicate a simultaneous substitution for many variables on a term  $e$  we will write e.g.,  $e[v_2 \mapsto t_2, v_1 \mapsto t_1]$ . To indicate a sequential substitution for many variables on a term  $e$  we will write  $e[v_1 \mapsto t_1] \circ [v_2 \mapsto t_2]$ , meaning  $(e[v_1 \mapsto t_1])[v_2 \mapsto t_2]$ .

An equation over *Sig* is a formula  $\forall \Gamma. t_1 = t_2 : s$  with  $t_1, t_2 \in \text{Terms}^s(\text{Sig}, \Gamma)$ ; we will omit the  $: s$  when doing so will not lead to confusion. A theory *Th* is a pair of a signature and a set of equations over that signature. In this paper, we will make use of a theory we call *Type*, which is displayed in Figure 4. Additional axioms, such as the associativity of  $+$ , can be added to *Type*, but doing so does not impact the examples in this paper.

Associated with a theory *Th* is a binary relation between (not necessarily ground) terms, called *provable equality*. We write  $Th \vdash \forall \Gamma. t = t' : s$  to indicate that the theory *Th* proves that terms  $t, t' \in \text{Terms}^s(\text{Sig}, \Gamma)$  are equal according to the usual rules of multi-sorted equational logic. From these rules it follows that provable equality is the smallest equivalence relation on terms that is a congruence, is closed under substitution, is closed under adding variables to contexts, and contains the equations of *Th*. Provable equality is semi-decidable in general (Bachmair *et al.*, 1989), but some theories are decidable. Formally,  $Th \vdash$  is defined by the inference rules in Figure 5.

A morphism of signatures  $F : \text{Sig}_1 \rightarrow \text{Sig}_2$  consists of

- a function  $F$  from sorts in  $\text{Sig}_1$  to sorts in  $\text{Sig}_2$ , and
- a function  $F$  from function symbols  $f : s_1 \times \dots \times s_n \rightarrow s$  in  $\text{Sig}_1$  to terms in

$$\text{Terms}^{F(s)}(\text{Sig}_2, \{v_1 : F(s_1), \dots, v_n : F(s_n)\})$$

To clearly indicate the context  $\{v_1, \dots, v_n\}$ , the function  $F(f)$  may be written in “ $\lambda$  notation,” i.e., as  $F(f) = \lambda v_1, \dots, v_n. g(v_1, \dots, v_n)$  for some term  $g$ , where the  $\lambda$  is omitted if  $n = 0$ .

For example, let  $Sig_1$  consist of two sorts,  $a, b$ , and one function symbol,  $f : a \rightarrow b$ , and let  $Sig_2$  consist of one sort,  $c$ , and one function symbol,  $g : c \rightarrow c$ . There are countably infinitely many morphisms  $F : Sig_1 \rightarrow Sig_2$ , one of which is defined as  $F(a) := c$ ,  $F(b) := c$ , and  $F(f) := \lambda v : c. g(g(v))$ . In the literature on algebraic specification, our definition of signature morphism is called a “derived signature morphism” (Mossakowski *et al.*, 2014).

The function  $F$  taking function symbols to terms can be extended to take terms to terms:

$$F(v) = v \quad F(f(t_1, \dots, t_n)) = F(f)[v_1 \mapsto F(t_1), \dots, v_n \mapsto F(t_n)]$$

As before, when we are defining the action of a specific  $F$  on a specific  $f : s_1 \times \dots \times s_n \rightarrow s$ , to make clear the variables we are using, we may write  $F(f) := \lambda v_1, \dots, v_n. \varphi$ , where  $\varphi$  may contain  $v_1, \dots, v_n$ . A *morphism of theories*  $F : Th_1 \rightarrow Th_2$  is a morphism of signatures that preserves provable equality of terms:

$$\begin{aligned} Th_1 \vdash \forall v_1 : s_1, \dots, v_n : s_n. t_1 = t_2 : s &\Rightarrow Th_2 \vdash \forall v_1 : F(s_1), \dots, v_n : F(s_n). F(t_1) \\ &= F(t_2) : F(s) \end{aligned}$$

In the theory *Type* (Figure 4), any permutation of  $A, B, \dots, Z$  induces a morphism  $Type \rightarrow Type$ , for example. Although morphisms of signatures are commonly used in the categorical approach to logic (Adámek *et al.*, 2011), such morphisms do not appear to be as commonly used in the traditional set-theoretic approach to logic. Checking that a morphism of signatures is a morphism of theories reduces to checking provable equality of terms and hence is semi-decidable.

**Remark.** Multi-sorted equational logic differs from single-sorted logic by allowing empty sorts (sorts that have no ground terms). Empty sorts are required by our formalism; without them, we could not express empty entities. For the theoretical development, this difference between multi-sorted and single-sorted logic can be safely ignored. But the fact that many algorithms are based on single-sorted logic means that care is required when implementing our formalism. For example, certain theorem proving methods based on Knuth–Bendix completion (Knuth & Bendix, 1970) require a ground term of every sort.

**Categorical remark.** From a theory  $Th$  we form a cartesian multi-category (Barr & Wells, 1995)  $\langle Th \rangle$  as follows. The objects of  $\langle Th \rangle$  are the sorts of  $Th$ . The elements of the hom-set  $s_1, \dots, s_k \rightarrow s$  of  $\langle Th \rangle$  are equivalence classes of terms of sort  $s$  in context  $\{v_1 : s_1, \dots, v_k : s_k\}$ , modulo the provable equality relation  $Th \vdash$ . Composition is defined by substitution. A morphism of theories  $F : Th_1 \rightarrow Th_2$  denotes a functor  $\langle F \rangle : \langle Th_1 \rangle \rightarrow \langle Th_2 \rangle$ . Although cartesian multi-categories are the most direct categorical semantics for theories, in many cases it is technically more convenient to work with product categories instead. Every cartesian multi-category generates a product category, and we often conflate the multi-category just described with the product category it generates, as is usually done in the categorical algebraic theories literature. For details, see Schultz *et al.* (2017).

$$\begin{aligned} \llbracket \text{Type} \rrbracket(\text{Nat}) &= \{0, 1, 2, \dots\} \\ \llbracket \text{Type} \rrbracket(\text{Char}) &= \{A, B, C, \dots\} \\ \llbracket \text{Type} \rrbracket(\text{String}) &= \{\text{nil}, A, B, \dots, AA, AB, \dots\} \\ \llbracket \text{Type} \rrbracket(\text{succ}) &= \{(0, 1), (1, 2), \dots\} \end{aligned}$$

Fig. 6. The term model  $\llbracket \text{Type} \rrbracket$  of theory  $\text{Type}$  (Figure 4).

### 3.2 Semantics

An algebra  $A$  over a signature  $\text{Sig}$  consists of

- a set of carriers  $A(s)$  for each sort  $s$ , and
- a function  $A(f) : A(s_1) \times \dots \times A(s_k) \rightarrow A(s)$  for each symbol  $f : s_1 \times \dots \times s_k \rightarrow s$ .

Let  $\Gamma := \{v_1 : s_1, \dots, v_n : s_n\}$  be a context. An  $A$ -environment  $\eta$  for  $\Gamma$  associates each  $v_i$  with an element of  $A(s_i)$ . The meaning of a term in  $\text{Terms}(\text{Sig}, \Gamma)$  relative to  $A$ -environment  $\eta$  for  $\Gamma$  is recursively defined as:

$$A\llbracket v \rrbracket\eta = \eta(v) \quad A\llbracket f(t_1, \dots, t_n) \rrbracket\eta = A(f)(A\llbracket t_1 \rrbracket\eta, \dots, A\llbracket t_n \rrbracket\eta)$$

An algebra  $A$  over a signature  $\text{Sig}$  is a *model* of a theory  $\text{Th}$  on  $\text{Sig}$  when  $\text{Th} \vdash \forall \Gamma. t = t' : s$  implies  $A\llbracket t \rrbracket\eta = A\llbracket t' \rrbracket\eta$  for all terms  $t, t' \in \text{Terms}^s(\text{Sig}, \Gamma)$  and  $A$ -environments  $\eta$  for  $\Gamma$ . Deduction in multi-sorted equational logic is sound and complete: two terms  $t, t'$  are provably equal in a theory  $\text{Th}$  if and only if  $t$  and  $t'$  denote the same element in every model of  $\text{Th}$ . One model of the theory  $\text{Type}$  (Figure 4) has carriers consisting of the natural numbers, the 26 character English alphabet, and all strings over the English alphabet. Another model of  $\text{Type}$  uses natural numbers modulo four as the carrier for  $\text{Nat}$ .

From a signature  $\text{Sig}$ , we form its *term algebra*  $\llbracket \text{Sig} \rrbracket$ , a process called *saturation*, as follows. The carrier set  $\llbracket \text{Sig} \rrbracket(s)$  is defined as the set of ground terms of sort  $s$ . The function  $\llbracket \text{Sig} \rrbracket(f)$  for  $f : s_1 \times \dots \times s_k \rightarrow s$  is defined as the function  $t_1, \dots, t_n \mapsto f(t_1, \dots, t_n)$ . From a theory  $\text{Th}$  on  $\text{Sig}$  we define its *term model*  $\llbracket \text{Th} \rrbracket$  to be the quotient of  $\llbracket \text{Sig} \rrbracket$  by the equivalence relation  $\text{Th} \vdash$ . In other words, the carrier set  $\llbracket \text{Th} \rrbracket(s)$  is defined as the set of equivalence classes of ground terms of sort  $s$  that are provably equal under  $\text{Th}$ . The function  $\llbracket \text{Th} \rrbracket(f)$  is  $\llbracket \text{Sig} \rrbracket(f)$  lifted to operate on equivalence classes of terms. To represent  $\llbracket \text{Th} \rrbracket$  on a computer, or to write down  $\llbracket \text{Th} \rrbracket$  succinctly, we must choose a *representative* for each equivalence class of terms; this detail can be safely ignored for now, but we will return to it in the implementation of our formalism (Section 5). When we must choose representatives for  $\llbracket \text{Th} \rrbracket$ , we will write  $nf_{\text{Th}}(e)$  to indicate the unique  $e' \in \llbracket \text{Th} \rrbracket$  such that  $\text{Th} \vdash e = e'$  (i.e., the *normal form* for  $e$  in  $\text{Th}$ ). For example, the term model of the algebraic theory  $\text{Type}$  (Figure 4) is shown in Figure 6.

A *morphism of algebras*  $h : A \rightarrow B$  on a signature  $\text{Sig}$  is a family of functions  $h(s) : A(s) \rightarrow B(s)$  indexed by sorts  $s$  such that

$$h(s)(A(f)(a_1, \dots, a_n)) = B(f)(h(s_1)(a_1), \dots, h(s_n)(a_n))$$

for every symbol  $f : s_1 \times \dots \times s_n \rightarrow s$  and  $a_i \in A(s_i)$ . We may abbreviate  $h(s)(a)$  as  $h(a)$  when  $s$  can be inferred. The term algebras for a signature  $\text{Sig}$  are initial among

all *Sig*-algebras: there is a unique morphism from the term algebra to any other *Sig*-algebra. Similarly, the term models are initial among all models. It is because of initiality that in many applications of equational logic to functional programming, such as algebraic datatypes (Mitchell, 1996), the intended meaning of a theory is its term model.

**Categorical remark.** Models of a theory  $Th$  correspond to functors  $Th \rightarrow \mathbf{Set}$ , and the term model construction yields an initial such model. That is, an algebraic theory  $Th$  denotes a cartesian multi-category,  $\langle Th \rangle$ , and the term model construction yields a functor  $\langle Th \rangle \rightarrow \mathbf{Set}$ . At the risk of confusion, we will also write  $\langle Th \rangle$  for the functor  $\langle Th \rangle \rightarrow \mathbf{Set}$ ; hence, we have  $\langle Th \rangle : \langle Th \rangle \rightarrow \mathbf{Set}$ . Morphisms between models correspond to natural transformations. A morphism of theories  $F : Th_1 \rightarrow Th_2$  induces a functor,  $\langle F \rangle : \langle Th_1 \rangle \rightarrow \langle Th_2 \rangle$  between the cartesian multi-categories  $\langle Th_1 \rangle$  and  $\langle Th_2 \rangle$ , as well as a natural transformation between the set-valued functors  $\langle Th_1 \rangle$  and  $\langle Th_2 \rangle$ .

#### 4 An equational formalism for functorial data migration

In this section, we describe how to use multi-sorted equational theories to define schemas and instances, and how to migrate data from one schema to another. To summarize, we proceed as follows; each of these steps is described in detail in this section. First, we fix an arbitrary multi-sorted equational theory  $Ty$  to serve as an ambient *type-side* or “background theory” against which we will develop our formalism. We say that the sorts in  $Ty$  are *types*. For example, we may define  $Ty$  to contain a sort  $\mathbf{Nat}$ ; function symbols  $0, 1, +, \times$ ; and equations such as  $0 + x = x$ . A *schema*  $S$  is an equational theory that extends  $Ty$  with new sorts (which we call *entities*), for example  $\mathbf{Person}$ ; unary function symbols between entities (which we call *foreign keys*) and from entities to types (which we call *attributes*), for example,  $\mathbf{father} : \mathbf{Person} \rightarrow \mathbf{Person}$  and  $\mathbf{age} : \mathbf{Person} \rightarrow \mathbf{Nat}$ ; and additional equations. An instance  $I$  on  $S$  is an equational theory that extends  $S$  with 0-ary constant symbols (which we call *generators*), such as  $\mathbf{Bill}$  and  $\mathbf{Bob}$ ; as well as additional equations, such as  $\mathbf{father}(\mathbf{Bill}) = \mathbf{Bob}$ . The intended meaning of  $I$  is its term model (i.e., the canonical model built from  $I$ -equivalence classes of terms). Morphisms of schemas and instances are defined as theory morphisms; i.e., mappings of sorts to sorts and function symbols to (open) terms that preserve entailment:  $h : C \rightarrow D$  exactly when  $C \vdash p = q$  implies  $D \vdash h(p) = h(q)$ . Our “uber-flower” query language is based on a generalization of for-let-where-return (flwr) notation, and generalizes the idea from relational database theory that conjunctive queries can be represented as “frozen” instances (Abiteboul *et al.*, 1995).

##### 4.1 Type-sides, schemas, instances, mappings, and transforms

Our formalism begins by fixing a specific multi-sorted equational theory  $Ty$ , which will be called the *type-side* of the formalism. The sorts in  $Ty$  are called *types*. The type-side is meant to represent the computational context within which our formalism will be deployed. For example, a type-side for SQL would contain sorts such as  $\mathbf{VARCHAR}$  and  $\mathbf{INTEGER}$  and functions such as  $\mathbf{LENGTH} : \mathbf{VARCHAR}$

→ INTEGER, as well as any user-defined scalar functions we wish to consider (e.g., squaring a number); a type-side for SK combinator calculus would contain a sort  $o$ , constants  $S, K : o$ , a function symbol  $\cdot : o \rightarrow o$ , and equations  $K \cdot x \cdot y = x$  and  $S \cdot x \cdot y \cdot z = x \cdot z \cdot (y \cdot z)$ . From a database point of view, choosing a particular multi-sorted equational theory  $Ty$  can be thought of as choosing the set of built-in and user-defined types and functions that can appear in schemas and queries.

Simply typed, first-order functional programs can be expressed using multi-sorted equational theories, and using a functional program as a type-side is a best-case scenario for the automated reasoning required to implement our formalism; see Section 5 for details. We will abbreviate “multi-sorted equational theory” as “theory” in this section.

A *schema*  $S$  on type-side  $Ty$  is a theory extending  $Ty$ . If  $s$  is a sort in  $S$  but not in  $Ty$ , we say that  $s$  is an *entity*; otherwise, that  $s$  is a *type*. (Note that although we say “entity,” the synonyms “entity type” and “entity set” are also common in database literature.)  $S$  must meet the following conditions:

1. If  $\forall \Gamma. t_1 = t_2$  is in  $S$  but not in  $Ty$ , then  $\Gamma = \{v : s\}$  where  $s$  is an entity.
2. If  $f : s_1 \times \dots \times s_n \rightarrow s$  is in  $S$  but not  $Ty$ , then  $n = 1$  and  $s_1$  is an entity. If  $s$  is an entity we say that  $f$  is a *foreign key*; otherwise, we say that  $f$  is an *attribute*.

In other words, every equation in a schema will have one of two forms:  $\forall v : t. v.p = v.p' : t'$ , where  $t$  and  $t'$  are entities and  $p$  and  $p'$  are sequences of foreign keys; or some combination of type-side functions applied to attributes, for example  $\forall v : t. v.p_1.att_1 + v.p_2.att_2 = v.att : t'$  where  $t$  is an entity and  $t'$  is a type. Due to these restrictions,  $S$  admits three sub-theories: the *type-side* of  $S$ , namely,  $Ty$ ; the *entity-side* of  $S$ , namely, the restriction of  $S$  to entities (written  $S_E$ ); and the *attribute-side* of  $S$ , namely, the restriction of  $S$  to attributes (written  $S_A$ ). We can also consider the entities and attributes together ( $S_{EA}$ ), and the attributes and type-side together ( $S_{AT}$ ). A morphism of schemas, or *schema mapping*,  $S_1 \rightarrow S_2$  on type-side  $Ty$  is a morphism of theories  $S_1 \rightarrow S_2$  that is the identity on  $Ty$ . An example schema *Emp* on type-side *Type* (Figure 4) is shown in Figure 7. We may draw the entity and attribute part of the schema in a graphical notation, with every sort represented as a dot, and the foreign keys and attributes represented as edges.

In schema *Emp* (Figure 7), *Emp* and *Dept* are sorts (entities) of employees and departments, respectively; *mgr* takes an employee to their manager; *secr* takes a department to its secretary; and *wrk* takes an employee to the department they work in. The equations are data integrity constraints saying that managers work in the same department as their employees, that secretaries work for the department they are the secretary for, and that the management hierarchy is two levels deep (this constraint ensures that  $\llbracket Emp_E \rrbracket$  is finite, a condition useful for our examples but not required by our implementation; see Section 5).

The first restriction on schemas (bullet 1 in the above list) rules out products of entities (for example, using a symbol *CommonBoss* :  $Emp \times Emp \rightarrow Emp$ ), and the second restriction on schemas (bullet 2) rules out the use of types as domains (for example, using a symbol *EmpNumber* :  $Nat \rightarrow Emp$ ). These restrictions are necessary to guarantee the existence of a right adjoint  $\Pi$  to  $\Delta$ , which we use to model product and filter operations.

$Sorts := \{Emp, Dept\}$   
 $Symbols := \{mgr : Emp \rightarrow Emp, wrk : Emp \rightarrow Dept, secr : Dept \rightarrow Emp\},$   
 $dname : Dept \rightarrow String, ename : Emp \rightarrow String$   
 $Equations := \{\forall v. v.mgr.wrk = v.wrk, \forall v. v = v.secr.wrk, \forall v. v.mgr.mgr = v.mgr\}$

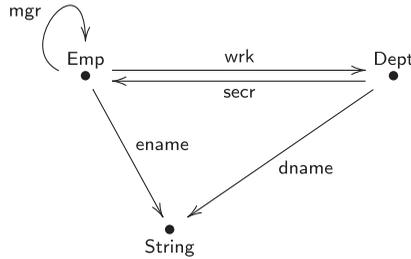


Fig. 7. Schema *Emp*, on type-side *Ty* (Figure 4).

An instance *I* on schema *S* is a theory extending *S*, meeting the following conditions:

1. If *s* is a sort in *I*, then *s* is a sort in *S*.
2. If  $\forall \Gamma. t_1 = t_2$  is in *I* but not *S*, then  $\Gamma = \emptyset$ .
3. If  $f : s_1 \times \dots \times s_n \rightarrow s$  is in *I* but not *S*, then  $n = 0$ . We say *f* is a generator of sort *s*.

That is, an instance only adds 0-ary symbols and ground equations to its schema. Mirroring a similar practice in database theory, we use the phrase *skolem term* to refer to a term in an instance that is not provably equal to a term entirely from the type-side, but whose sort is a type. Although skolem terms are very natural in database theory, skolem terms wreak havoc in the theory of algebraic datatypes, where their existence typically implies an error in a datatype specification that causes computation to get “stuck” (Mitchell, 1996).

Similarly to how schemas admit sub-theories for entity, attribute, and type-sides, an instance *I* contains sub-theories for entities (*I<sub>E</sub>*), attributes (*I<sub>A</sub>*), and types (*I<sub>T</sub>*). Note that *I<sub>T</sub>* may not be the ambient type-side *Ty*, because *I* can declare new constant symbols whose sorts are types (so-called skolem variables), as well as additional equations; for example,  $infinity : Nat$  and  $succ(infinity) = infinity$ . A morphism of instances, or *transform*,  $h : I_1 \rightarrow I_2$  is a morphism of theories  $I_1 \rightarrow I_2$  that is the identity on *S*, and the requirement of identity on *S* rules out the “attribute problem” from Figure 3.

The intended meaning of an instance *I* is its term model,  $\llbracket I \rrbracket$ . In practice, the term model  $\llbracket I \rrbracket$  will often have an infinite type-side, but  $\llbracket I_{EA} \rrbracket$  will be finite. Therefore, our implementation computes  $\llbracket I_{EA} \rrbracket$ , as well as an instance *talg(I)* called the *type-algebra*<sup>1</sup> for *I*. The type-algebra *talg(I)* is an instance on the empty schema

<sup>1</sup> Technically, it is  $\llbracket talg(I) \rrbracket$  that is a *Ty*-algebra, and *talg(I)* presents this algebra. But we will almost never be interested in  $\llbracket talg(I) \rrbracket$ , so to save space we will refer to the equational theory *talg(I)* as *I*’s type-algebra.

$$\text{Symbols} := \{a, b, c : \text{Emp}, m, s : \text{Dept}\}$$

$$\text{Eqs} := \{a.\text{ename} = \text{Al}, c.\text{ename} = \text{Carl}, m.\text{dname} = \text{Math}, \\ a.\text{wrk} = m, b.\text{wrk} = m, s.\text{secr} = c, m.\text{secr} = b\}$$

			Emp			
			ID	ename	mgr	wrk
Dept			a	Al	a.mgr	m
ID	dname	secr	b	b.ename	b.mgr	m
m	Math	b	c	Carl	c.mgr	s
s	s.dname	c	a.mgr	a.mgr.ename	a.mgr	m
			b.mgr	b.mgr.ename	b.mgr	m
			c.mgr	c.mgr.ename	c.mgr	s

The type-algebra extends  $Type$  with:

$$\text{Symbols} := \{m.\text{dname}, s.\text{dname}, a.\text{ename}, b.\text{ename}, c.\text{ename}, \\ a.\text{mgr.ename}, b.\text{mgr.ename}, c.\text{mgr.ename} : \text{String}\}$$

$$\text{Eqs} := \{a.\text{ename} = \text{Al}, c.\text{ename} = \text{Carl}, m.\text{dname} = \text{Math}\}$$

Fig. 8. Instance  $Inst$  on schema  $Emp$  (Figure 7).

over  $I_T$ . For every attribute  $att : s \rightarrow t$  in  $S$ , and every term  $e \in \llbracket I_E \rrbracket(s)$ , the type-algebra contains a generator  $e.att : t$ . We call these generators *observables* because they correspond to type-valued observations one can make about an entity. Observables have the form  $e.fk_1 \dots fk_n.att$ ; i.e., have a 0-ary constant symbol as a head, followed by a possibly empty list of foreign keys, followed by an attribute. We define the function  $trans : \text{Terms}^t(I, \emptyset) \rightarrow \text{talg}(I)$ , for every type (non-entity) sort  $t$ , as

$$trans(e.fk_1 \dots fk_n.att) := nf_I(e.fk_1 \dots fk_n).att \text{ for observables}$$

$$trans(f(e_1, \dots, e_n)) := f(trans(e_1), \dots, trans(e_n)) \text{ otherwise}$$

By  $nf_I(x)$ , we mean the normal form for  $x$  in  $\llbracket I \rrbracket$ ; see Section 3.2. The equations for  $\text{talg}(I)$  are the images under  $trans$  of the (necessarily ground) equations of  $I$  but not  $S$  and all substitution instances of the equations at types in  $S$  but not  $Ty$ . Note that  $\text{talg}(I)$  does not present  $\llbracket I_T \rrbracket$  (the restriction of  $I$  to types), rather,  $\text{talg}(I)$  presents  $\llbracket I \rrbracket_T$  (the skolem terms of  $\llbracket I \rrbracket$  and their relationships).

We visually present term models using a set of tables, with one table per entity, with an ID column corresponding to the carrier set. Sometimes, we will present the type-algebra as well. An instance on the  $Emp$  schema and its denotation are shown in Figure 8.

In many cases, we would like for an instance  $I$  to be a *conservative extension* of its schema  $S$ , meaning that for all terms  $t, t' \in \text{Terms}^s(S, \Gamma)$ ,  $I \vdash \forall \Gamma. t = t' : s$  if and only if  $S \vdash \forall \Gamma. t = t' : s$ . (Similarly, we may also want schemas to conservatively extend their type-sides.) For example,  $\text{Emp} \not\vdash \text{Al} = \text{Carl} : \text{String}$ , but there is an  $\text{Emp}$ -instance  $I$  for which  $I \vdash \text{Al} = \text{Carl} : \text{String}$ . In the context of “deductive databases” (Abiteboul *et al.*, 1995) (databases that are represented intensionally, as theories, rather than extensionally, as tables) such as our formalism, non-conservativity is usually regarded as non-desireable (Ghilardi *et al.*, 2006), although nothing in our formalism requires conservativity. Checking for conservativity is decidable for the description logic underlying OWL (Ghilardi *et al.*, 2006), but not decidable for multi-sorted equational logic (and hence our formalism), and not decidable for the formalism of EDs (Fagin *et al.*, 2005b) that underlies much work on relational data integration (the chase fails when conservativity is violated). In Section 5.3, we give a simple algorithm that soundly approximates conservativity. Note that the  $\Delta$  and  $\Pi$  migration functors preserve the conservative extension property, but  $\Sigma$  does not; hence, one may want to be careful when using  $\Sigma$ . (More pedantically,  $\Delta$  and  $\Pi$  preserve type-algebras, but  $\Sigma$  does not.)

**Remark.** There is a precise sense in which our definition of transform corresponds to the definition of database homomorphism in relational database theory. Recall that in database theory (Abiteboul *et al.*, 1995), a schema is a triple  $(\text{dom}, \text{null}, R)$ , where  $\text{dom}$  is a set (called the *domain*),  $\text{null}$  is a set disjoint from  $\text{dom}$  (called the *labeled nulls*), and  $R$  is a set of relation names and arities; an instance  $I$  is a set of relations over  $\text{dom} \cup \text{null}$ , indexed by  $R$ , of appropriate arity; and a homomorphism  $h : I_1 \rightarrow I_2$  is a function  $h : \text{dom} \cup \text{null} \rightarrow \text{dom} \cup \text{null}$  that is constant on  $\text{dom}$  and such that  $(c_1, \dots, c_n) \in I_1(R)$  implies  $(h(c_1), \dots, h(c_n)) \in I_2(R)$  for every  $R$  of arity  $n$ . If we interpret a term model  $\llbracket I \rrbracket$  as a relational instance by considering every skolem term in  $\llbracket I \rrbracket$  to be a labeled null and every non-skolem term to be a domain value, then a transform of instances in our formalism induces a homomorphism of the encoded relational instances. In this encoding,  $\text{dom}$  is playing the role of a free (equation-less), discrete (function-less) type-side.

**Categorical remark.** By forgetting the entity/attribute distinction, we can consider a schema  $S$  as a single algebraic theory,  $\hat{S}$ ; the cartesian multi-category  $\langle \hat{S} \rangle$  is called the *collage* of  $S$ . A schema mapping  $F : S \rightarrow T$  is then a functor between collages  $\langle \hat{S} \rangle \rightarrow \langle \hat{T} \rangle$  that is the identity on  $Ty$ . More pedantically, a schema  $S$  is a profunctor  $\langle S \rangle : \langle S_E \rangle^{op} \times \langle Ty \rangle \rightarrow \mathbf{Set}$  that preserves products in  $\langle Ty \rangle$ . The observables from an entity  $e \in S_E$  to type  $ty \in Ty$  are given by  $\langle S \rangle(e, ty)$ . A schema mapping  $F : S \rightarrow T$  denotes a functor  $\langle F_E \rangle : \langle S_E \rangle \rightarrow \langle T_E \rangle$  and a natural transformation  $\langle S \rangle \Rightarrow \langle T \rangle \circ (\langle F_E \rangle^{op} \times id)$ :

$$\begin{array}{ccc}
 \langle S_E \rangle^{op} \times \langle Ty \rangle & \xrightarrow{\langle S \rangle} & \mathbf{Set} \\
 \langle F_E \rangle^{op} \times id \downarrow & \searrow \langle T \rangle & \uparrow \\
 \langle T_E \rangle^{op} \times \langle Ty \rangle & & 
 \end{array}$$

#### 4.2 Functorial data migration

We are now in a position to define the data migration functors. We first fix a type-side (multi-sorted equational theory),  $Ty$ . The following are proved in Schultz *et al.* (2017):

- The schemas on  $Ty$  and their mappings form a category.
- The instances on a schema  $S$  and their transforms form a category,  $S\text{-Inst}$ .
- The models of  $S$  and their homomorphisms obtained by applying  $\llbracket \cdot \rrbracket$  to  $S\text{-Inst}$  form a category,  $\llbracket S\text{-Inst} \rrbracket$ , which is equivalent to, but not equal to,  $S\text{-Inst}$ .
- A schema mapping  $F : S \rightarrow T$  induces a unique functor  $\Sigma_F : S\text{-Inst} \rightarrow T\text{-Inst}$  defined by substitution,  $\Sigma_F(I) := F(I)$ , with a right adjoint,  $\Delta_F : T\text{-Inst} \rightarrow S\text{-Inst}$ , which itself has a right adjoint,  $\Pi_F : S\text{-Inst} \rightarrow T\text{-Inst}$ .
- A schema mapping  $F : S \rightarrow T$  induces a unique functor  $\llbracket \Delta_F \rrbracket : \llbracket T\text{-Inst} \rrbracket \rightarrow \llbracket S\text{-Inst} \rrbracket$  defined by composition,  $\llbracket \Delta_F \rrbracket(I) := I \circ F$ , with a left adjoint,  $\llbracket \Sigma_F \rrbracket : \llbracket S\text{-Inst} \rrbracket \rightarrow \llbracket T\text{-Inst} \rrbracket$ , and a right adjoint  $\llbracket \Pi_F \rrbracket : \llbracket S\text{-Inst} \rrbracket \rightarrow \llbracket T\text{-Inst} \rrbracket$ .

Although  $\Sigma_F$  and  $\llbracket \Delta_F \rrbracket$  are canonically defined, their adjoints are only defined up to unique isomorphism. The canonically defined migration functors enjoy properties that the other data migration functors do not, such as  $\Sigma_F(\Sigma_G(I)) = \Sigma_{F \circ G}(I)$  and  $\llbracket \Delta_F \rrbracket(\llbracket \Delta_G \rrbracket(I)) = \llbracket \Delta_{F \circ G} \rrbracket(I)$  (for the other functors, these are not equalities, but unique isomorphisms).

It is possible to give explicit formulae to define the three data migration functors  $\Delta, \Sigma, \Pi$  (Schultz *et al.*, 2017). However, we have found that it is more convenient to work with two derived data migration functors,  $\Delta \circ \Pi$  and  $\Delta \circ \Sigma$ , which we describe in the next section. Therefore, we now simply describe examples of  $\Delta, \Sigma, \Pi$  in Figures 2, 9, and 10. Because these examples display instances as tables, rather than equational theories, we are actually illustrating  $\llbracket \Delta \rrbracket, \llbracket \Sigma \rrbracket, \llbracket \Pi \rrbracket$ .

Figures 2 and 9 shows a schema mapping  $F$  that takes two distinct source entities,  $N1$  and  $N2$ , to the target entity  $N$ . The  $\llbracket \Delta_F \rrbracket$  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 between  $N1$  and  $N2$ , the  $\llbracket \Delta_F \rrbracket$  functor populates it so that  $N$  can be recovered by joining  $N1$  and  $N2$ . The  $\llbracket \Pi_F \rrbracket$  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  $\llbracket \Sigma_F \rrbracket$  functor disjointly unions  $N1$  and  $N2$ ; because  $N1$  and  $N2$  are not union compatible (have different columns),  $\llbracket \Sigma_F \rrbracket$  creates null values. When there is a foreign key between  $N1$  and  $N2$ ,  $\llbracket \Sigma_F \rrbracket$  merges the tuples that are related by the foreign key (Roberson & Wyss, 2004), resulting in a join. As these two examples illustrate,  $\Delta$  can be thought of as “foreign-key aware” projection,  $\Pi$  can be thought of as a product followed by a filter (which can result in a join), and  $\Sigma$  can be thought of as a (not necessarily union compatible) disjoint union followed by a merge (which can also result in a join).

Figure 10 shows a traditional “data exchange setting” (Fagin *et al.*, 2005b); data on a source schema about amphibians must be migrated onto a target schema about animals, where the target schema contains a data integrity constraint enforcing that

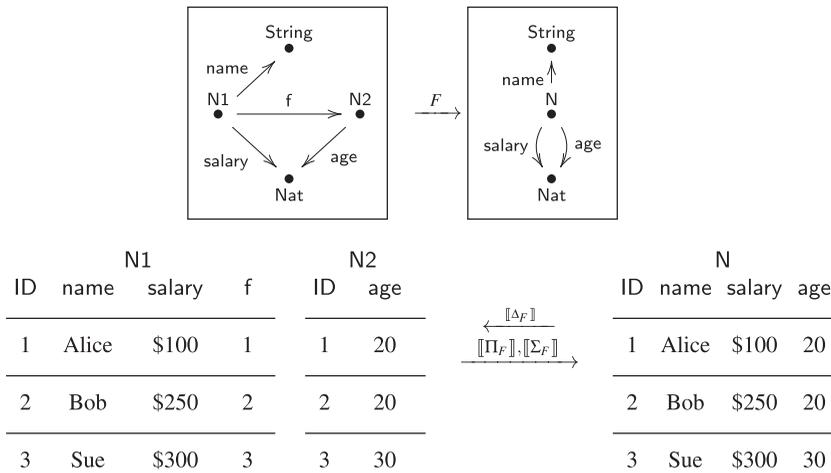


Fig. 9. Example functorial data migrations, with foreign keys.

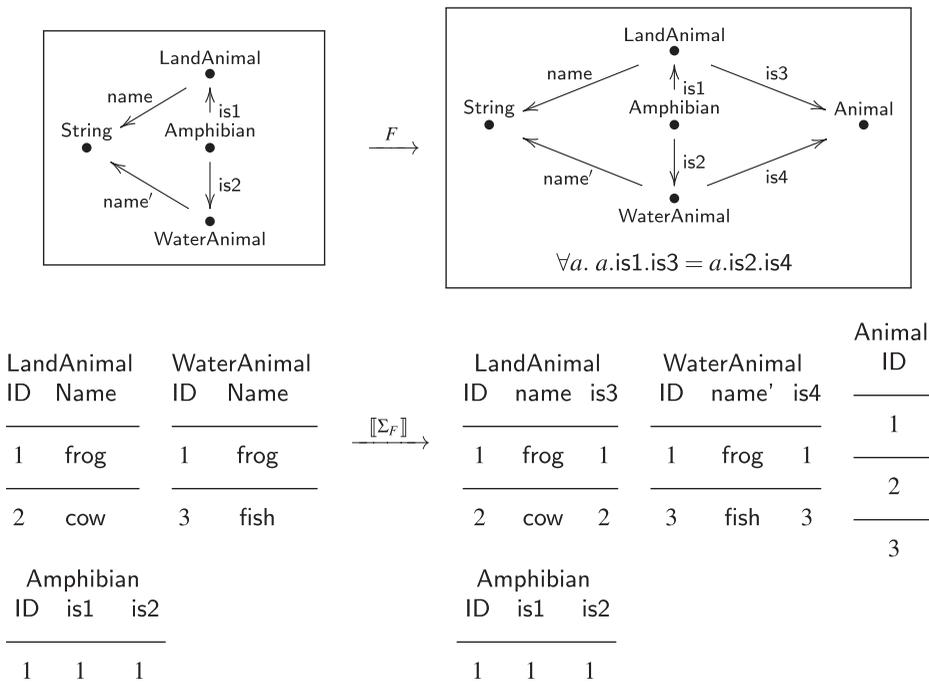


Fig. 10. An example  $\Sigma$  data migration, with path equalities.

each amphibian is only counted as a single animal. The schema mapping  $F$  is an inclusion, and  $\llbracket \Sigma_F \rrbracket$  has precisely the desired semantics.

### 4.3 Uber-flower queries

It is possible to form a query language directly from schema mappings. This is the approach we took in Spivak & Wisnesky (2015), where a query is defined to be a triple of schema mappings  $(F, G, H)$  denoting  $\llbracket \Sigma_F \rrbracket \circ \llbracket \Pi_G \rrbracket \circ \llbracket \Delta_H \rrbracket$ . Suitable conditions on

$F, G, H$  guarantee closure under composition, computability using relational algebra, and other properties desirable in a query language. In practice, however, we found this query language to be challenging to program. Having to specify entire schema mappings is onerous; it is difficult to know how to use the data migration functors to accomplish any particular task without a thorough understanding of category theory; and as a kind of “join all,”  $\Pi$  is expensive to compute. Hence, in Schultz *et al.* (2017), we developed a new syntax, which we call *uber-flower* syntax because it generalizes *flwr* (for-let-where-return) syntax (a.k.a. select-from-where syntax, a.k.a. comprehension syntax; Grust, 2004). We have found *uber-flower* syntax to be more concise, easier to program, and easier to implement than the language based on triples of schema mappings in Spivak & Wisnesky (2015).

An *uber-flower*  $Q : S \rightarrow T$ , where  $S$  and  $T$  are schemas on the same type-side, induces a data migration  $eval(Q) : S\text{-Inst} \rightarrow T\text{-Inst} \cong \Delta_G \circ \Pi_F$  and an adjoint data migration  $coeval(Q) : T\text{-Inst} \rightarrow S\text{-Inst} \cong \Delta_F \circ \Sigma_G$  for some  $X$ ,  $F : S \rightarrow X$ ,  $G : T \rightarrow X$ . In fact, all data migrations of the form  $\Delta \circ \Pi$  can be expressed as the *eval* of an *uber-flower*, and all migrations of the form  $\Delta \circ \Sigma$  can be expressed as the *coeval* of an *uber-flower*. In Sections 4.3.4 and 4.3.5, we describe the correspondence between *uber-flowers* and data migration functors in detail. In the remainder of this section we describe *uber-flowers*, but defer a description of how to (co-)evaluate them to Sections 5.4 and 5.5.

A *tableau* (Abiteboul *et al.*, 1995) over a schema  $S$  is a pair of

- a context over  $S$ , called the *for* clause,  $fr$ , and
- a set of quantifier-free equations between terms in  $Terms(S, fr)$ , called the *where* clause  $wh$ .

Associated with a *tableau* over  $S$  is a canonical  $S$ -instance, the “frozen” instance (Abiteboul *et al.*, 1995). In our formalism, a *tableau* trivially becomes an instance by the validity-preserving *Herbrandization* process (the dual of the satisfiability-preserving *Skolemization* process) that “freezes” variables into fresh constant symbols. For example, we can consider the *tableau*  $(\{v_1 : Emp, v_2 : Dept\}, v_1.wrk = v_2, v_1.ename = Peter)$  to be an *Emp*-instance with generators  $v_1, v_2$ . In this paper, we may silently pass between a *tableau* and its frozen instance.

An *uber-flower*  $S \rightarrow T$  consists of, for each entity  $t \in T$ :

- a *tableau*  $(fr_t, wh_t)$  over  $S$  and,
- for each attribute  $att : t \rightarrow t' \in T$ , a term  $[att]$  in  $Terms^t(S, fr_t)$ , called the *return* clause for  $att$ , and
- for each foreign key  $fk : t \rightarrow t' \in T$ , a transform  $[fk]$  from the *tableau* for  $t'$  to the *tableau* for  $t$  (note the reversed direction), called the *keys* clause for  $fk$ ,
- such that an equality-preservation condition holds. We defer a description of this condition until Section 4.3.1.

We prefer to use  $fr_t, wh_t, [att], [fk]$  notation when discussing a single *uber-flower*. When we are discussing many queries  $Q_1, \dots, Q_n$ , we will write  $Q_k(t), Q_k(att), Q_k(fk)$  to indicate  $(fr_t, wh_t), [att]$ , and  $[fk]$ , respectively, for  $Q_k$ .

We usually require that the *for* clauses in an *uber-flower* only bind variables to entities, not to types (e.g.,  $v : Person$  is allowed, but  $v : Nat$  is not). While not strictly

$$\begin{array}{l}
 \top, \perp : \text{Bool} \quad \neg : \text{Bool} \rightarrow \text{Bool} \quad \wedge, \vee : \text{Bool} \times \text{Bool} \rightarrow \text{Bool} \\
 x \vee y = y \vee x \qquad x \wedge y = y \wedge x \\
 x \vee (y \vee z) = (x \vee y) \vee z \qquad x \wedge (y \wedge z) = (x \wedge y) \wedge z \\
 x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z) \qquad x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z) \\
 x \vee \neg x = \top \qquad x \wedge \neg x = \perp \\
 x \vee \perp = x \qquad x \wedge \top = x
 \end{array}$$

Fig. 11. An equational axiomatization of Boolean algebra.

necessary, there are two reasons for preferring this restriction. First, in practice, types will almost always be infinite, so the data migrations induced by a non-restricted uber-flower would often return infinite instances. Second, the restriction ensures that the induced data migrations are *domain independent* (Abiteboul *et al.*, 1995), allowing some evaluations of uber-flowers to be computed using relational algebra (Spivak & Wisnesky, 2015). Semantically, this restriction means that evaluations of uber-flowers correspond to migrations of the form  $\Delta_G \circ \Pi_F$ , where  $F$  is surjective on attributes, a condition described in Section 4.3.4.

In Figure 12 we present an uber-flower, from our *Emp* schema to our *Emp* schema (Figure 7), which when evaluated makes each employee their own boss and appends their old boss’s name to their name. Note that Dept is copied over unchanged; only Emp changes.

**Remark.** Because Boolean algebra can be equationally axiomatized, evaluation of uber-flowers can express queries that might not be considered conjunctive in certain relational settings. For example, when our type-side contains Boolean algebra (Figure 11), evaluation of uber-flowers can express queries such as  $Q(R) := \{x \in R \mid P(x) \vee \neg P'(x) = \top\}$ . In addition, instances can contain skolem terms of type Bool, implying the existence of truth values besides  $\top, \perp$ . Similarly, evaluation of uber-flowers can express non-computable queries whenever a type-side contains an equational theory that is Turing-complete. In such cases, the type-side will not have a decidable equality relation and the theorem proving methods employed by the implementation (Section 5) will diverge or fail.

**Categorical remark.** An uber-flower is syntax for a structure that has several equivalent formulations. One is induced by a cospan of schemas (Section 4.3.5). Another is a certain kind of profunctor between schemas: let  $Q : S \rightarrow T$  be an uber-flower on type-side  $Ty$ . Then,  $Q$  denotes a *bimodule* (Schultz *et al.*, 2017), i.e., a functor  $\langle Q \rangle : \langle \hat{T} \rangle^{op} \rightarrow \langle S\text{-Inst} \rangle$  where  $\langle Q \rangle(t) = y(t)$  for all  $t \in \langle Ty \rangle$ , where  $y : \langle \hat{S} \rangle^{op} \rightarrow \langle S\text{-Inst} \rangle$  is the Yoneda embedding.

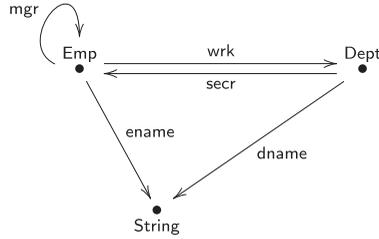
### 4.3.1 Verification conditions for uber-flower well-formedness

Let  $Q : S \rightarrow T$  be an uber-flower. To verify that  $Q$  is well-formed, for every equation in  $T$  we must verify an induced set of equations on  $S$ . There are two kinds of equations in  $T$  that we must consider, both of which require the notion of *path* to describe. A *path* in  $T$ , namely  $p : t_0 \rightarrow t_n := fk_1 \dots fk_n$ , is a sequence of foreign keys,  $fk_n : t_{n-1} \rightarrow t_n$ , and we write  $[p]$  to indicate the substitution  $[fk_1] \circ \dots \circ [fk_n]$  taking each  $v : s \in fr_{t_n}$  to some term in  $Terms^s(S, fr_{t_0})$  (when  $n = 0$  this is the identity substitution on  $fr_{t_0}$ ). One kind of equation to verify is an equality between

```

Dept = for d:Dept
  where
  return (dname -> d.dname)
  keys (secr -> [e -> d.secr])

Emp = for e:Emp
  where
  return (ename -> e.ename + e.mgr.ename)
  keys (mgr -> [e -> e], wrk -> [d -> e.wrk])
    
```



Verification conditions:

$$\begin{aligned}
 \forall v. v.mgr.wrk = v.wrk : Dept &\rightsquigarrow Emp \cup e : Emp \vdash d[d \mapsto e.wrk][e \mapsto e] = d[d \mapsto e.wrk] : Dept \\
 \forall v : Dept. v = v.secr.wrk : Dept &\rightsquigarrow Emp \cup d : Dept \vdash d = d[d \mapsto e.wrk][e \mapsto d.secr] : Dept \\
 \forall v. v.mgr.mgr = v.mgr : Emp &\rightsquigarrow Emp \cup e : Emp \vdash e[e \mapsto e][e \mapsto e] = e[e \mapsto e] : Emp
 \end{aligned}$$

Fig. 12. The uber-flower  $Promote : Emp \rightarrow Emp$ .

entities of paths with a shared head variable in a singleton entity context:

$$\forall v : t. v.p = v.p' : t'$$

For each variable  $u : s \in fr_{t'}$ , we must verify the following:

$$S \cup (fr_t, wh_t) \vdash u[p] = u[p'] : t'$$

The other kind of equation to verify is an equality between types of arbitrary terms where paths share head variables in a singleton entity context, for example,

$$\forall v : t. c + v.p'.att' = v.p.att : t'$$

Here, we must check, for example,

$$S \cup (fr_t, wh_t) \vdash c + [att'][p'] = [att][p] : t'$$

Figure 12 shows an uber-flower and its verification conditions. In that figure  $\varphi \rightsquigarrow \psi$  means that equation  $\varphi \in T$  generates verification condition  $\psi$ . Rather than simply give the verification conditions, the figure illustrates how the verification conditions are obtained. For example, the first verification condition

$$\begin{aligned}
 \forall v. v.mgr.wrk = v.wrk : Dept &\rightsquigarrow Emp \cup e : Emp \vdash d[d \mapsto e.wrk][e \mapsto e] = d[d \\
 &\mapsto e.wrk] : Dept
 \end{aligned}$$

means that the equation  $eq := \forall v. v.mgr.wrk = v.wrk : Dept$  induces the (tautological) verification condition  $Emp \cup e : Emp \vdash e.wrk = e.wrk : Dept$  by starting with a substitution instance  $eq' := d.mgr.wrk = d.wrk : Dept$  of  $eq$ , and then

applying the keys clauses (substitutions) of the uber-flower,  $[wrk] = [d \mapsto e.wrk]$  and  $[mgr] = [e \mapsto e]$ , to  $eq'$ . The notation  $Emp \cup$  means that we can use the equations from the  $Emp$  schema (Figure 7) to prove the desired equality.

From a practical standpoint, well-formed queries are guaranteed to only materialize instances that obey their data integrity constraints, so runtime checking of data integrity constraints are not needed. Mathematically, a query must be well-formed to even be considered a query.

### 4.3.2 Morphisms of uber-flowers

Let  $Q_1, Q_2 : S \rightarrow T$  be uber-flowers. For every foreign key  $fk : t \rightarrow t'$ , we have transforms  $Q_1(f) : fr_t^1 \rightarrow fr_{t'}^1$ ,  $Q_2(f) : fr_t^2 \rightarrow fr_{t'}^2$ . A morphism  $h : Q_1 \rightarrow Q_2$  is, for each entity  $t \in T$ , a transform of frozen instances  $h(t) : fr_t^1 \rightarrow fr_t^2$ , such that for every foreign key  $fk : t \rightarrow t' \in T$ , and every  $v' \in fr_{t'}^1$ ,

$$S \cup wh_t^2 \vdash v'[h(t)'] [Q_2(f)] = v' [Q_1(f)] [h(t)]$$

and for every attribute  $att : t \rightarrow t' \in T$ ,

$$S \cup wh_t^2 \vdash Q_2(att) = Q_1(att)[h(t)]$$

The morphism  $h$  induces a transform  $eval(h) : eval(Q_2)(I) \rightarrow eval(Q_1)(I)$  for each  $S$ -instance  $I$  (in this way,  $evals$  of uber-flowers are similar to relational conjunctive queries, for which  $Q \rightarrow Q'$  implies  $\forall J, Q'(J) \subseteq Q(J)$ ) and a transform  $coeval(h) : coeval(Q_1)(J) \rightarrow coeval(Q_2)(J)$  for each  $T$ -instance  $J$ , described in Sections 5.4.2 and 5.5.2, respectively.

### 4.3.3 Composing uber-flowers

Let  $Q_1 : A \rightarrow B$  and  $Q_2 : B \rightarrow C$  be uber-flowers on the same type-side. We define the uber-flower  $Q_2 \circ Q_1 : A \rightarrow C$  as follows. For every entity  $c \in C$ ,

- We define (the frozen  $A$ -instance of) the for and where clause at  $c$  of  $Q_2 \circ Q_1$  to be the frozen  $A$ -instance that is obtained by applying  $coeval(Q_1) : B\text{-Inst} \rightarrow A\text{-Inst}$  to the frozen  $B$ -instance for  $c$  in  $Q_2$ :

$$(Q_2 \circ Q_1)(c) := coeval(Q_1)(Q_2(c))$$

- Similarly, the transform associated with a foreign key  $fk : c \rightarrow c' \in C$  is

$$(Q_2 \circ Q_1)(fk) := coeval(Q_1)(Q_2(fk))$$

- To define the term associated with an attribute  $att : c \rightarrow c' \in C$ , we first define the instance  $y(c')$  (a so-called “representable” Barr & Wells (1995) instance) to be the  $B$ -instance with a single generator  $v : c'$ . We then define a transform/substitution from  $y(c')$  to the frozen  $B$ -instance for  $c$  in  $Q_2$ , namely,  $h : y(c') \rightarrow Q_2(c) := v \mapsto Q_2(att)$ . Finally, we define

$$(Q_2 \circ Q_1)(att) := coeval(Q_1)(h(v))$$

### 4.3.4 Converting data migrations to uber-flowers

Let  $F : S \rightarrow T$  be a schema mapping on a type-side  $Ty$ . In this section we define two uber-flowers,  $Q_F : S \rightarrow T$  and  $Q^F : T \rightarrow S$  such that

$$eval(Q^F) \cong \Delta_F \quad coeval(Q^F) \cong \Sigma_F \quad eval(Q_F) \cong \Pi_F \quad coeval(Q_F) \cong \Delta_F$$

We now describe  $Q^F$  and  $Q_F$ :

- The for clause for  $Q^F : T \rightarrow S$  at entity  $s \in S$  is defined to have a single variable,  $v_s : F(s)$ , and  $Q^F$  has an empty where clause. For each foreign key,  $fk : s \rightarrow s' \in S$ ,  $F(fk)$  is ( $\alpha$ -equivalent to) a term in  $Terms^{F(s')}(T, \{v_s : F(s)\})$  and we define  $Q^F(fk) : Q^F(s') \rightarrow Q^F(s)$  to be the transform  $v_{s'} \mapsto F(fk)$ . For each attribute  $att : s \rightarrow s' \in S$ ,  $s'$  is a type and  $F(att)$  is ( $\alpha$ -equivalent to) a term in  $Terms^{s'}(T, \{v_s : F(s)\})$  and we define  $Q^F(att)$  to be  $F(att)$ . For example, the  $Q^F$  that corresponds to Figure 9 is

```
N1 := for vN1 : N,
      return name -> vN1.name, salary -> vN1.salary
      keys f -> [vN2 -> vN1.f]

N2 := for vN2 : N,
      return age -> vN2.age
```

- The frozen instance (for/where clause) for  $Q_F : S \rightarrow T$  at entity  $t \in T$  is defined to be  $\Delta_F(y(t))$ , where  $y(t)$  is the instance with a single generator  $\{v_t : t\}$ . For each foreign key  $fk : t \rightarrow t' \in T$ , we define the transform  $Q_F(fk) : Q_F(t') \rightarrow Q_F(t)$  to be  $\Delta_F(v_{t'} \mapsto v_t.fk)$ . For each attribute  $att : t \rightarrow t' \in T$ ,  $v_t.att \in Terms^{t'}(T, \{v_t : t\})$  and  $trans(v_t.att) \in talg(\{v_t : t\})$ . Since  $\Delta$  preserves type algebras, we have  $trans(v_t.att) \in talg(\Delta_F(\{v_t : t\}))$ , and hence we can define  $Q_F(att)$  to be  $trans(v_t.att)$ . For example, the  $Q_F$  that corresponds to Figure 9 is (note that we write “ $x$ ” to indicate an  $x \in talg(y(N))$ ):

```
y(N) = vN : N
N := for vN1 : N1, vN2 : N2
     where vN1.f = vN2, vN1.name='vN.name',
           vN1.salary = 'vN.salary', vN2.salary = 'vN.salary'
     return name -> vN1.name, salary -> vN2.salary, age -> vN2.age
```

For  $Q_F$  to be an uber-flower that obeys the restriction that variables in for clauses only bind entities and not types,  $F$  must be *surjective on attributes*. This semi-decidable condition implies that for every  $I$ , there will be no skolem terms in  $talg(\Delta_F(I))$ , i.e.,  $talg(\Delta_F(I)) \cong Ty$ . Formally,  $F$  is surjective on attributes when for every attribute  $att : t \rightarrow t' \in T$ , there exists an entity  $s \in S$  such that  $F(s) = t$  and there exists an  $e \in Terms^{t'}(T, \{v : s\})$  such that  $T \vdash \forall v : F(s). F(e) = v.att$ .

### 4.3.5 Converting uber-flowers to data migrations

An uber-flower  $Q : S \rightarrow T$ , where  $S$  and  $T$  are schemas on the same type-side  $Ty$ , induces a data migration  $eval(Q) : S\text{-Inst} \rightarrow T\text{-Inst} \cong \Delta_G \circ \Pi_F$  and adjoint data

migration  $coeval(Q) : T\text{-Inst} \rightarrow S\text{-Inst} \cong \Delta_F \circ \Sigma_G$  for some  $X, F : S \rightarrow X, G : T \rightarrow X$ . In this section, we construct  $X, F$ , and  $G$ . First, we define a schema  $X$  such that  $S \subseteq X$  and we define  $F : S \hookrightarrow X$  to be the inclusion mapping. We start with

$$En(X) := En(S) \sqcup En(T) \quad Att(X) := Att(S) \quad Fk(X) \subseteq Fk(S) \sqcup Fk(T)$$

Then, for each entity  $t \in T$  and each  $v : s$  in  $fr_t$  (the frozen instance for  $t$  in  $Q$ ), we add a foreign key to  $X$ :

$$(v, s, t) : t \rightarrow s \in Fk(X)$$

Let us write  $\sigma_x$  for the substitution  $[v_k \mapsto x.(v_k, s_k, t), \forall v_k : s_k \in fr_t]$ . For each equation  $e = e' \in wh_t$  we add an equation to  $X$ :

$$\forall x : t. e\sigma_x = e'\sigma_x \in Eq(X)$$

and for each foreign key  $fk : t \rightarrow t'$  and for each  $v' : s' \in fr_{t'}$ , we add an equation to  $X$ :

$$\forall x : t. x.fk.(v', s', t') = v'[fk]\sigma_x \in Eq(X) \tag{1}$$

This almost completes the schema  $X$ , but we will need to add the equations of  $T$ , suitably translated, to  $X$ . To do so, we must first define  $G : T \rightarrow X$  to be the identity on entities and foreign keys, and on attributes, we define

$$G(att : t \rightarrow t') := \forall x : t. [att]\sigma_x$$

Finally, to complete  $X$  we add the images of  $T$ 's equations under  $G$  to  $X$ .

For example, the schema  $X$  for the uber-flower *Promote* (Figure 12) is shown in Figure 13. Rather than simply give the equations of the schema  $X$ , the figure illustrates how the equations conditions are obtained. For example, the first equation

$$\forall x. x.wrk.d = d[d \mapsto e.wrk][e \mapsto x.e] \equiv x.e.wrk$$

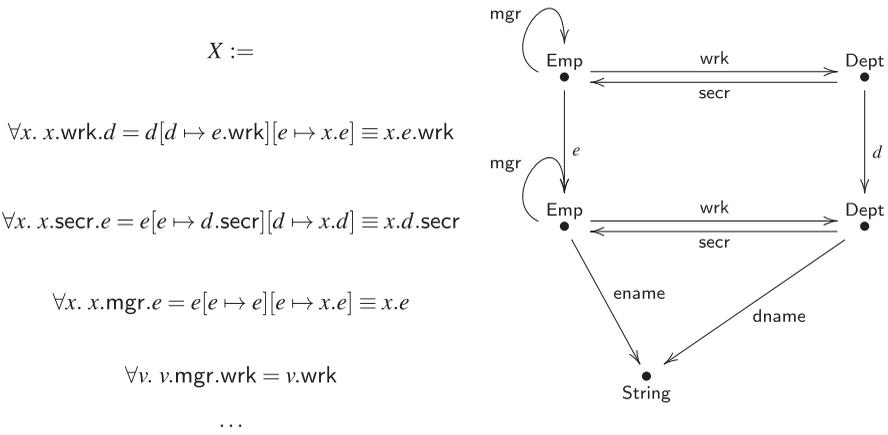
means that schema  $X$  contains the equation  $\forall x. x.wrk.d = x.e.wrk$ , which was obtained from foreign key  $wrk : Emp \rightarrow Dept$  and for-bound variable  $d : Dept$  by equation 1.

#### 4.3.6 ED syntax for uber-flowers

Intriguingly, the intermediate schema and schema mappings that are created when translating uber-flowers into data migrations, as described in the previous section, suggest an alternative syntax for uber-flowers that resembles the syntax of second-order EDs (Fagin *et al.*, 2005a). The uber-flower *Promote* (Figure 12) is shown as a data migration in Figure 13, and we can express the intermediate schema and mapping in Figure 13 using the following second-order ED:

$$\exists e : Emp_{src} \rightarrow Emp_{dst}, \exists d : Dept_{src} \rightarrow Dept_{dst}.$$

$$\begin{aligned} \forall x. d(wrk_{src}(x)) &= wrk_{dst}(e(x)) \wedge \forall x. e(secr_{src}(x)) \\ &= secr_{dst}(d(x)) \wedge \forall x. e(mgr_{src}(x)) = e(x) \end{aligned}$$



$$G(ename) := \forall x. [ename][e \mapsto x.e] \equiv x.e.ename + x.e.mgr.ename$$

Fig. 13. Uber-flower *Promote* (Figure 12) as a data migration.

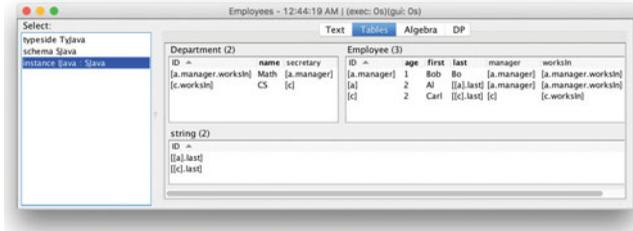


Fig. 14. The AQL tool displaying an instance.

We do not understand how our formalism relates to second-order EDs, but the implementation of our formalism in the AQL tool allows users to input uber-flowers using the above second-order ED syntax.

### 5 Implementation: the AQL tool

We have implemented our formalism in the open-source AQL tool, which can be downloaded at <http://categoricaldata.net/aql.html>. In this section, we discuss certain implementation issues that arise in negotiating between syntax and semantics, and provide algorithms for key parts of the implementation: deciding equality in equational theories, saturating theories into term models, checking conservativity of equational theories, (co-)evaluating queries, and (co-)pivoting (converting instances into schemas). Even though the goal of the AQL tool is merely to prove that algebraic data integration can be done, we close with a discussion about the AQL tool’s performance.

#### 5.1 Deciding equality in equational theories

Many constructions involving equational theories, including uber-flower (co-) evaluation, depend on having a decision procedure for provable equality in the

Equations	Rewrite Rules	
$1 * x = x$	$1 * x \rightsquigarrow x$	$x * 1 \rightsquigarrow x$
$x^{-1} * x = 1$	$x^{-1} * x \rightsquigarrow 1$	$(x^{-1})^{-1} \rightsquigarrow x$
$(x * y) * z = x * (y * z)$	$(x * y) * z \rightsquigarrow x * (y * z)$	$x * x^{-1} \rightsquigarrow 1$
	$x^{-1} * (x * y) \rightsquigarrow y$	$x * (x^{-1} * y) \rightsquigarrow y$
	$1^{-1} \rightsquigarrow 1$	$(x * y)^{-1} \rightsquigarrow y^{-1} * x^{-1}$

Fig. 15. Knuth–Bendix completion for group theory.

theory. A decidable equational theory is said to have a decidable *word problem*. The word problem is obviously semi-decidable: to prove if two terms (words)  $p$  and  $q$  are equal under equations  $E$ , we can systematically enumerate all of the (usually infinite) consequences of  $E$  until we find  $p = q$ . However, if  $p$  and  $q$  are not equal, then this enumeration will never stop. In practice, not only is enumeration computationally infeasible, but for uber-flower (co-)evaluation, we require a true decision procedure: an algorithm that, when given  $p$  and  $q$  as input, will always terminate with “equal” or “not equal.” Hence, we must look to efficient, but incomplete, automated theorem proving techniques to decide word problems.

The AQL tool provides a built-in theorem prover based on Knuth–Bendix completion (Knuth & Bendix, 1970): from a set of equations  $E$ , it attempts to construct a system of rewrite rules (oriented equations),  $R$ , such that  $p$  and  $q$  are equal under  $E$  if and only if  $p$  and  $q$  rewrite to syntactically equal terms (so-called *normal forms*) under  $R$ . We demonstrate this with an example. Consider the equational theory of groups, on the left, in Figure 15. Knuth–Bendix completion yields the rewrite system on the right in Figure 15. To see how these rewrite rules are used to decide the word problem, consider the two terms  $(a^{-1} * a) * (b * b^{-1})$  and  $b * ((a * b)^{-1} * a)$ . Both of these terms rewrite to 1 under the above rewrite rules; hence, we conclude that they are provably equal. In contrast, the two terms  $1 * (a * b)$  and  $b * (1 * a)$  rewrite to  $a * b$  and  $b * a$ , respectively, which are not syntactically the same; hence, we conclude that they are not provably equal.

The details of how the Knuth–Bendix algorithm works are beyond the scope of this paper. However, we make several remarks. First, Knuth and Bendix’s original algorithm (Knuth & Bendix, 1970) can fail even when a rewrite system to decide a word problem exists; for this reason, we use the more modern, “unfailing” variant of Knuth–Bendix completion (Bachmair *et al.*, 1989). Second, first-order, simply-typed functional programs are equational theories that are already complete in the sense of Knuth–Bendix. Third, specialized Knuth–Bendix algorithms exist for particular kinds of theories; a particular algorithm (Kapur & Narendran, 1985) for theories where all function symbols are 0-ary or unary, such as for the entity and attribute parts of our schemas, works well in practice.

## 5.2 Saturating theories into term models

Many constructions involving equational theories, including uber-flower (co-)evaluation, depend on having a procedure, called *saturation*, for constructing finite

term models from theories. This process is semi-computable: there are algorithms that will construct a finite term model if it exists, but diverge if no finite term model exists. The AQL tool has two different methods for saturating theories: theories where all function symbols are 0-ary or unary can be saturated using an algorithm for computing Left-Kan extensions (Bush *et al.*, 2003), and arbitrary theories can be saturated by using a decision procedure for the theory's word problem as follows. Let  $Th$  be an equational theory, and define the *size* of a term in  $Th$  to be the height of the term's abstract syntax tree; for example,  $\max(x.sal, x.mgr.sal)$  has size of three. We construct  $\llbracket Th \rrbracket$  in stages: first, we find all not provably equal terms of size 0 in  $Th$ ; call this  $\llbracket Th \rrbracket^0$ . Then, we add to  $\llbracket Th \rrbracket^0$  all not provably equal terms of size 1 that are not provably equal to a term in  $\llbracket Th \rrbracket^0$ ; call this  $\llbracket Th \rrbracket^1$ . We iterate this procedure, obtaining a sequence  $\llbracket Th \rrbracket^0, \llbracket Th \rrbracket^1, \dots$ . If  $\llbracket Th \rrbracket$  is indeed finite, then there will exist some  $n$  such that  $\llbracket Th \rrbracket^n = \llbracket Th \rrbracket^{n+1} = \llbracket Th \rrbracket$  and we can stop. Otherwise, our attempt to construct  $\llbracket Th \rrbracket$  will run forever: it is not decidable whether a given theory  $Th$  has a finite term model.

Note that the model  $\llbracket Th \rrbracket$  computed using the above procedure is technically not the canonical term model for the theory; rather, we have constructed a model that is isomorphic to the canonical term model by choosing representatives for equivalence classes of terms under the provable equality relation. Depending on how we enumerate terms, we can end up with different models.

Saturation is used for constructing tables from instances to display to the user, and for (co-)evaluating queries on instances. In general, the type-side  $Ty$  of an instance  $I$  will be infinite, so we cannot saturate the equational theory of the instance directly (i.e.,  $\llbracket I \rrbracket$  is often infinite). For example, if the type-side of  $I$  is the free group on one generator  $a$ , then  $\llbracket I \rrbracket$  will contain  $a$ ,  $a * a$ ,  $a * a * a$ , and so on. Hence, as described in Section 4, the AQL tool computes the term model for only the entity and attribute part of  $I$  (namely,  $\llbracket I_{EA} \rrbracket$ ), along with an instance (equational theory) called the type-algebra of  $I$  (namely,  $talg(I)$ ). The pair  $(\llbracket I_{EA} \rrbracket, talg(I))$  is sufficient for all of AQL's purposes.

The AQL tool supports an experimental feature that we call "computational type-sides." The mathematics behind this feature have not been fully worked out, but it provides a mechanism to connect AQL to other programming languages. An  $\mathcal{L}$ -valued model of  $talg(I)$  is similar to a (set-valued) model of  $talg(I)$ , except that instead of providing a carrier set for each sort in  $talg(I)$ , an  $\mathcal{L}$ -valued model provides a type in  $\mathcal{L}$ , and instead of providing a function for each symbol in  $talg(I)$ , an  $\mathcal{L}$ -valued model provides an expression in  $\mathcal{L}$ . For example, if  $\mathcal{L} = \text{Java}$ , then we can interpret `String` as `java.lang.String`, `Nat` as `java.lang.Integer`, `+ : String  $\rightarrow$  String` as `java.lang.String.append`, etc. (Note that our  $\mathcal{L}$ -models are on  $talg(I)$ , not  $Ty$ ; so an  $\mathcal{L}$ -model must provide a meaning for the skolem terms in  $talg(I)$ , which can be tricky.) Given an  $\mathcal{L}$ -model  $M$ , we can take the image of  $\llbracket I_{EA} \rrbracket$  under  $M$  by replacing each term  $talg(t) \ni t \in \llbracket I_{EA} \rrbracket$  with the value of  $t$  in  $M$ . We write this as  $M(\llbracket I_{EA} \rrbracket)$ . The pair  $(M(\llbracket I_{EA} \rrbracket), M)$  can be used by the AQL tool in many situations where  $(\llbracket I_{EA} \rrbracket, talg(A))$  is expected, for example, displaying instances (see Figure 16), and (co-)evaluating uber-flowers. Formalizing computational type-sides is an important area for future work.

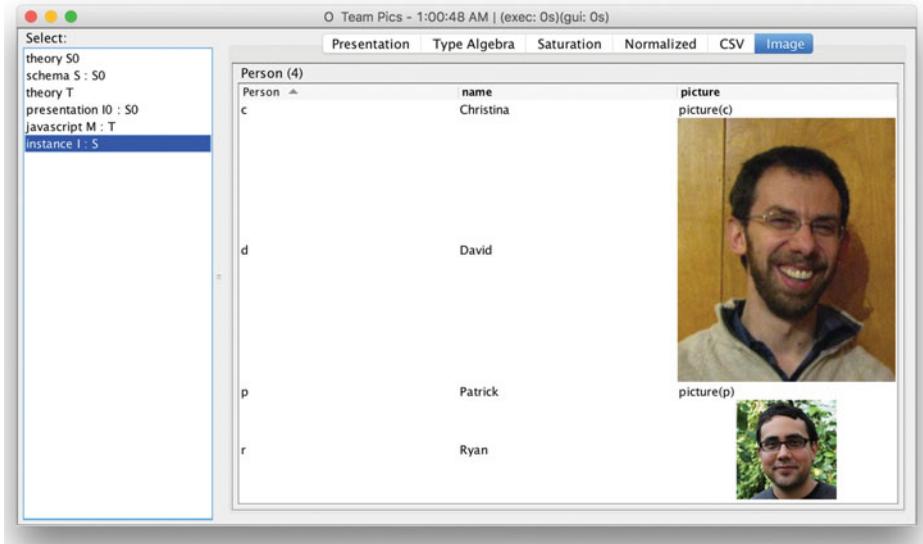


Fig. 16. The AQL tool displaying an instance with a computational type-side.

### 5.3 Deciding that a theory conservatively Extends another

As described in Section 4, we may want instances to conservatively extend their schemas, where a theory  $Th_2$  conservatively extends  $Th_1$  when  $Th_1 \vdash \forall \Gamma. t = t' : s$  iff  $Th_2 \vdash \forall \Gamma. t = t' : s$  for all  $t, t' \in Terms^s(Th_1, \Gamma)$  for every  $\Gamma$ . Conservativity in equational logic is not decidable, and the only system we are aware of that automates conservativity checks in a language at least as expressive as equational logic is CCC (Lüth *et al.*, 2005). In this section, we give a simple algorithm that soundly but incompletely checks that a theory  $Th_2$  conservatively extends a theory  $Th_1$  by showing that  $Th_2$  freely extends  $Th_1$ .

Let  $Th_1$  be an equational theory, and let  $Th_2$  extend  $Th_1$  with new sorts, symbols, and equations. We will simplify the presentation of  $Th_2$  by repeatedly looking for equations of the form  $g = t$ , where  $g$  is a generator (0-ary symbol) of  $Th_2$  but not of  $Th_1$ , and  $t$  does not contain  $g$ ; we then substitute  $g \mapsto t$  in  $Th_2$ . If after no more substitutions are possible, all equations in  $Th_2$  are either reflexive or provable in  $Th_1$ , then  $Th_2$  is conservative (actually, free) over  $Th_1$ . For example, we can show that the theory

$$\{\text{infinity} : \text{Nat}, \text{undef} : \text{Nat}, \text{infinity} = 0, \text{undef} = 1, \text{infinity} = \text{undef}\}$$

is not conservative over *Type* (Figure 4), because the simplification process yields the non-reflexive equation  $0 = 1$ , which is not provable in *Type*. However, the algorithm is far from complete. The theory

$$\{+ : \text{Nat} \times \text{Nat} \rightarrow \text{Nat}, \text{infinity} : \text{Nat}, \text{undef} : \text{Nat}, \text{infinity}+1 = \text{undef}+1\}$$

does not pass our check, even though it is a conservative extension of *Type*. Developing a better conservativity checker is an important area for future work, lest we inadvertently “damage our ontologies” (Ghilardi *et al.*, 2006). The process of

repeatedly substituting  $g \mapsto t$ , where  $g$  is a generator in an instance's type-algebra and  $t$  is a type-side term is also used by the AQL tool to simplify the display of tables by biasing the tables to display e.g., 45 instead of e.g. age.bill when 45 and age.bill are provably equal.

#### 5.4 Evaluating uber-flowers

Although it is possible to evaluate an uber-flower by translation into a data migration of the form  $\Delta \circ \Pi$ , we have found that in practice it is faster to evaluate such queries directly, using algorithms that extend existing join algorithms from relational query processing. In this section, we describe such an algorithm that intuitively extends the most basic join algorithm, “nested loops join” (Garcia-Molina *et al.*, 2008).

Let  $Q : S \rightarrow T$  be an uber-flower and let  $I$  be an  $S$ -instance. We now describe how to compute the instance (theory)  $eval(Q)(I)$ . First, we copy the generators and equations of the type-algebra  $talg(I)$  into  $eval(Q)(I)$ . Then, for every target entity  $t \in T$ , we perform the following:

- We define the generators of entity  $t$  in  $eval(Q)(I)$  to be those  $\llbracket I_{EA} \rrbracket$  environments for  $fr_t$  that satisfy  $wh_t$ . Formally, we represent these environments as ground substitutions  $fr_t \rightarrow Terms(I, \emptyset)$  and define, where  $fr_t := \{\overline{v_i} : \overline{s_i}\}$ :

$$eval(Q)(I)(t) := \{\overline{[v_i \mapsto \overline{e_i}]} \mid I \vdash eq[\overline{v_i \mapsto \overline{e_i}}], \forall eq \in wh_t, \forall e_i \in \llbracket I_{EA} \rrbracket(s_i)\}$$

- For each attribute  $att : t \rightarrow t' \in T'$ , we have a term  $[att] \in Terms'(S, fr_t)$  from the return clause for  $t$ . For every substitution  $\sigma \in eval(Q)(t)$ , we have  $[att]\sigma \in Terms'(S, \emptyset)$ , and we add

$$\sigma.att = trans([att]\sigma) \in eval(Q)(I)$$

The reason that  $trans([att]\sigma) \in Terms'(eval(Q)(I), \emptyset)$  is because  $trans([att]\sigma) \in Terms'(talg(I), \emptyset)$  and  $talg(I) \subseteq eval(Q)(I)$ . See Section 4.1 for the definition of  $trans$ .

- For each foreign key  $fk : t \rightarrow t' \in T$ , we have a transform from the frozen instance for  $t'$  to the frozen instance for  $t$  from the keys clause for  $t$ , which can be thought of as a substitution  $[fk] : fr_{t'} \rightarrow Terms(S, fr_t)$ . For every substitution  $\sigma : fr_t \rightarrow Terms(S, \emptyset) \in eval(Q)(t)$ , we add the equation:

$$\sigma.fk = \sigma \circ [fk] \in eval(Q)(I)$$

We know that  $\sigma \circ [fk] \in eval(Q)(I)(t')$  because  $[fk]$  is a transform, not an arbitrary substitution.

Note that in order to build the instance  $eval(Q)(I)$ , we have effectively constructed the term model  $\llbracket eval(Q)(I)_{EA} \rrbracket$ , and then “de-saturated” it into an equational theory. The description above is implemented in the AQL tool by a simple nested loop join algorithm: for each target entity  $t$ , to find those  $\llbracket I_{EA} \rrbracket$  environments for  $fr_t := \{\overline{v_i} : \overline{s_i}\}$  satisfying  $wh_t$ , compute a temporary table  $\tau := \Pi_i \llbracket I_{EA} \rrbracket(s_i)$  and then filter  $\tau$  by  $wh_t$ , using provable equality in  $I$ . The AQL tool contains additional implementations of

query evaluation based on more sophisticated algorithms such as hash-join (Garcia-Molina *et al.*, 2008), but we do not describe these algorithms here.

To make the above description concrete, we will now evaluate the uber-flower  $Promote : Emp \rightarrow Emp$  from Figure 12 on the instance  $Inst$  from Figure 8, which in turn is on the schema  $Emp$  from Figure 7 on the type-side  $Type$  from Figure 4. Our goal is to compute the instance (equational theory)  $eval(Promote, Inst)$ . We start by copying  $talg(Inst)$  into  $eval(Promote, Inst)$ . Next, we process the tableau. We start with target entity  $Dept \in T$ . The from and where clauses give us a set of substitutions  $\{[d \mapsto m], [d \mapsto s]\}$ , which are the generators of  $eval(Promote, Inst)$  at entity  $Dept$ . The return clause adds equations  $[d \mapsto m].dname = m.dname$  and  $[d \mapsto s].dname = s.dname$ ; note that  $s.dname$  is one of the generators from  $talg(Inst)$ , and  $s$  will not be a term in  $eval(Promote, Inst)$ . The keys clause for  $secl : Dept \rightarrow Emp$  adds equations  $[d \mapsto m].secl = [e \mapsto b]$  and  $[d \mapsto s].secl = [e \mapsto c]$ ; we have not added  $[e \mapsto b]$  and  $[e \mapsto c]$  to  $eval(Promote, Inst)$  yet but we will momentarily. Note that so far, we have simply copied the table  $Dept$  from  $Inst$  to  $eval(Promote, Inst)$ , up to isomorphism. We next consider the target entity  $Emp \in T$ . The from and where clause give us a set of substitutions  $\{[e \mapsto a], [e \mapsto b], [e \mapsto c], [e \mapsto a.mgr], [e \mapsto b.mgr], [e \mapsto c.mgr]\}$ , which are the generators of  $eval(Promote, Inst)$  at entity  $Emp$ . The return clause adds equations such as  $[e \mapsto a].ename = a.ename + a.mgr.ename$ , where  $a.ename$  and  $a.mgr.ename$  are generators in  $talg(Inst)$ . The keys clause for  $mgr : Emp \rightarrow Emp$  adds equations such as  $[e \mapsto a].mgr = [e \mapsto a]$ , and the keys clause for  $wrk : Emp \rightarrow Dept$  adds equations such as  $[e \mapsto a].wrk = [d \mapsto m]$  (we added  $[d \mapsto m]$  to  $eval(Promote, Inst)$  when processing the target entity  $Emp$ ). The entire instance is displayed in Figure 17.

### 5.4.1 Evaluating uber-flowers on transforms

Let  $Q : S \rightarrow T$  be an uber-flower and let  $h : I \rightarrow J$  be a transform of  $S$ -instances  $I$  and  $J$ . Our goal is to define the transform  $eval(Q)(h) : eval(Q)(I) \rightarrow eval(Q)(J)$ . For each target entity  $t \in T$ , consider the generators in  $eval(Q)(I)(t)$ : they will be the ground substitutions  $fr_t \rightarrow Terms(S, \emptyset)$  satisfying  $wh_t$ . We will map each such substitution to a substitution (generator) in  $eval(Q)(J)(t)$ . Given such a substitution  $\sigma := [v_0 : s_0 \mapsto e_0, \dots, v_n : s_n \mapsto e_n]$ , where  $e_i \in Terms^{si}(S, \emptyset)$ , we define

$$eval(Q)(\sigma) := [v_0 : s_0 \mapsto nf_{J_E}(h(e_1)), \dots, v_n : s_n \mapsto nf_{J_E}(h(e_n))]$$

In general,  $h(e_i)$  need not appear in  $[[J_E]]$ , so we must use  $nf$  (Section 3.2) to find the normal form of  $h(e_i)$  in  $J_E$ .

### 5.4.2 Evaluating morphisms of uber-flowers

If  $Q_1, Q_2 : S \rightarrow T$  are uber-flowers, a morphism  $h : Q_1 \rightarrow Q_2$  is, for each entity  $t \in T$ , a morphism from the frozen instance for  $t$  in  $Q_1$  to the frozen instance for  $t$  in  $Q_2$ , and it induces a transform  $eval(h) : eval(Q_2)(I) \rightarrow eval(Q_1)(I)$  for every  $S$ -instance  $I$ ; we now show how to compute  $eval(h)$ . Let  $t \in T$  be an entity and  $fr_t^1 := \{v_1, \dots, v_n\}$  be the for clause for  $t$  in  $Q_1$ . The generators of  $eval(Q_2)(I)$  are

Generators := {m.dname, s.dname, a.ename, b.ename, c.ename : String,  
a.mgr.ename, b.mgr.ename, c.mgr.ename : String,  
 $[d \mapsto m], [d \mapsto s] : \text{Dept},$   
 $[e \mapsto a], [e \mapsto b], [e \mapsto c], [e \mapsto a.\text{mgr}], [e \mapsto b.\text{mgr}], [e \mapsto c.\text{mgr}] : \text{Emp}$ }

Eqs := {a.ename = Al, c.ename = Carl, m.dname = Math,  
 $[d \mapsto m].\text{secr} = [e \mapsto b], [d \mapsto s].\text{secr} = [e \mapsto c], [d \mapsto m].\text{dname} = \text{Math}, \dots$ }

			Emp			
			ID	ename	mgr	wrk
Dept			$[e \mapsto a]$	Al + a.mgr.ename	$[e \mapsto a]$	$[d \mapsto m]$
ID	dname	secr	$[e \mapsto b]$	b.ename + b.mgr.ename	$[e \mapsto b]$	$[d \mapsto m]$
$[d \mapsto m]$	Math	$[e \mapsto b]$	$[e \mapsto c]$	Carl + c.mgr.ename	$[e \mapsto c]$	$[d \mapsto s]$
$[d \mapsto s]$	s.dname	$[e \mapsto c]$	$[e \mapsto a.\text{mgr}]$	a.mgr.ename + a.mgr.ename	$[e \mapsto a.\text{mgr}]$	$[d \mapsto m]$
			$[e \mapsto b.\text{mgr}]$	b.mgr.ename + b.mgr.ename	$[e \mapsto b.\text{mgr}]$	$[d \mapsto m]$
			$[e \mapsto c.\text{mgr}]$	c.mgr.ename + c.mgr.ename	$[e \mapsto c.\text{mgr}]$	$[d \mapsto s]$

Fig. 17. Evaluation of uber-flower *Promote* (Figure 12) on *Inst* (Figure 8).

substitutions  $\sigma : fr_t^2 \rightarrow \llbracket I_E \rrbracket$ , and

$$eval(h)(\sigma) := [v_1 \mapsto nf_{I_E}(h(v_1)\sigma), \dots, v_n \mapsto nf_{I_E}(h(v_n)\sigma)]$$

In the above, we must use *nf* (Section 3.2) to find appropriate normal forms.

### 5.5 Co-evaluating uber-flowers

Although it is possible to co-evaluate an uber-flower by translation into a data migration of the form  $\Delta \circ \Sigma$ , we have implemented co-evaluation directly. Let  $Q : S \rightarrow T$  be an uber-flower and let  $J$  be a  $T$ -instance. We are not aware of any algorithm in relational database theory that is similar to *coeval*( $Q$ ); intuitively, *coeval*( $Q$ )( $J$ ) products the frozen instances of  $Q$  with the input instance  $J$  and equates the resulting pairs based on either the frozen part or the input part. We now describe how to compute the  $S$ -instance (theory) *coeval*( $Q$ )( $J$ ). First, we copy the generators and equations of the type-algebra *talg*( $J$ ) into *coeval*( $Q$ )( $J$ ). We define *coeval*( $Q$ )( $I$ ) to be the smallest theory such that, for every target entity  $t \in T$ , where  $fr_t := \{v_1 : s_1, \dots, v_n : s_n\}$ ,

- $\forall (v : s) \in fr_t$ , and  $\forall j \in \llbracket J \rrbracket(t)$ ,

$$(v, j) : s \in coeval(Q)(J)$$

- $\forall e = e' \in wh_t$ , and  $\forall j \in \llbracket J \rrbracket(t)$ ,

$$(e = e')[v_1 \mapsto (v_1, j), \dots, v_n \mapsto (v_n, j)] \in coeval(Q)(J)$$

- $\forall att : t \rightarrow t' \in T$ , and  $\forall j \in \llbracket J \rrbracket(t)$ ,

$$trans(\llbracket J \rrbracket(att)(j)) = [att][v_1 \mapsto (v_1, j), \dots, v_n \mapsto (v_n, j)] \in coeval(Q)(J)$$

recall that  $[att] \in Terms^t(S, fr_t)$  is the return clause for attribute  $att$  and  $trans : Terms^t(J, \emptyset) \rightarrow talg(J)$  is defined in Section 4.1.

- $\forall fk : t \rightarrow t' \in T$ , and  $\forall j \in \llbracket J \rrbracket(t)$ , and  $\forall (v' : s') \in fr_{t'}$ ,

$$(v', \llbracket J \rrbracket(fk)(j)) = v'[fk][v_1 \mapsto (v_1, j), \dots, v_n \mapsto (v_n, j)] \in coeval(Q)(J)$$

recall that the substitution  $[fk] : fr_{t'} \rightarrow Terms(S, fr_t)$  is the keys clause for  $fk$ .

The co-evaluation of the uber-flower  $Promote : Emp \rightarrow Emp$  from Figure 12 on the instance  $Inst$  from Figure 8 is in fact isomorphic to the evaluation of  $promote$  (Figure 17); the reason is that evaluation and co-evaluation of  $Promote$  are semantically both projections ( $\Delta$ -only operations).

### 5.5.1 Co-evaluating uber-flowers on transforms

Let  $Q : S \rightarrow T$  be an uber-flower and let  $h : I \rightarrow J$  be a transform of  $T$ -instances  $I$  and  $J$ . Our goal is to define the transform  $coeval(Q)(h) : coeval(Q)(I) \rightarrow coeval(Q)(J)$ . For each entity  $t \in T$ , and for every  $(v : s) \in fr_t$ , and for every  $j \in \llbracket J_E \rrbracket(t)$ , we define

$$coeval(Q)((v, j)) := (v, nf_{J_E}(h(j)))$$

As was the case for evaluation of uber-flowers on transforms, we must use  $nf$  (Section 3.2) to find appropriate normal forms.

### 5.5.2 Co-evaluating morphisms of uber-flowers

If  $Q_1, Q_2 : S \rightarrow T$  are uber-flowers, a morphism  $h : Q_1 \rightarrow Q_2$  is, for each entity  $t \in T$ , a morphism from the frozen instance for  $t$  in  $Q_1$  to the frozen instance for  $t$  in  $Q_2$ , and it induces a transform  $coeval(h) : coeval(Q_1)(J) \rightarrow coeval(Q_2)(J)$  for every  $T$ -instance  $J$ ; in this section, we show how to compute  $coeval(h)$ . Let  $t \in T$  be an entity. The generators of  $coeval(Q_1)(J)$  are pairs  $(v, j)$  with  $v : s \in fr_t^{Q_1}$  and  $j \in \llbracket J_E \rrbracket(t)$ . Define

$$coeval(h)((v, j)) := (v', j).fk_1 \dots fk_n \quad \text{where } h(v) := v'.fk_1 \dots fk_n$$

## 5.6 The unit and co-unit of the co-eval $\dashv$ eval adjunction

Let  $Q : S \rightarrow T$  be an uber-flower. Then  $coeval(Q)$  is left adjoint to  $eval(Q)$ , i.e.,  $coeval(Q) \dashv eval(Q)$ . This means that the set of morphisms  $coeval(Q)(I) \rightarrow J$  is isomorphic to the set of morphisms  $I \rightarrow eval(Q)(J)$  for every  $I, J$ . The unit and

co-unit of the adjunction, defined here, describe this isomorphism. Let  $I$  be a  $S$ -instance. The component at  $I$  of the co-unit transform  $\epsilon_I : coeval(Q)(eval(Q)(I)) \rightarrow I$  is defined as

$$\epsilon_I((v_k, [v_1 \mapsto e_1, \dots, v_n \mapsto e_2])) := e_k, \forall k \in \{1, \dots, n\}$$

Let  $J$  be a  $T$ -instance. The component at  $I$  of the unit transform  $\eta_J : J \rightarrow eval(Q)(coeval(Q), J)$  is defined as

$$\eta_J(j) := [v_1 \mapsto (v_1, nf_{J_E}(j)), \dots, v_n \mapsto (v_n, nf_{J_E}(j))]$$

### 5.7 Converting instances to schemas

Let  $S$  be a schema on type-side  $Ty$  and let  $I$  be an  $S$ -instance. We now describe how to convert  $I$  into a schema, written  $\int I$ , an operation we call “pivoting.” Then, we describe an alternative way to convert  $I$  into a schema, written  $\oint I$ , an operation we call “co-pivoting.” We indicate (co-)pivots using integral signs because related categorical constructions, such as the Grothendieck construction or co-ends (Barr & Wells, 1995) are often indicated using integral signs; however, here we are *defining* the schemas  $\int I$  and  $\oint I$ , and our integral notation is not meant to indicate any existing construction.

The *pivot* of  $I$  is defined to be a schema  $\int S$  on  $Ty$ , a mapping  $F : \int I \rightarrow S$ , and a  $\int I$ -instance  $J$ , such that  $\Sigma_F(J) = I$ , defined as follows. First, we copy  $talg(I)$  into  $J$ . Then, for every entity  $s \in S$ , and every  $i \in \llbracket I_E \rrbracket(s)$ , we define

$$i : s \in \int I \quad F(i) := s \quad i : s \in J$$

and for every attribute  $att : s \rightarrow s' \in S$ ,

$$(i, att) : i \rightarrow s' \in \int I \quad F((i, att)) := att \quad i.(i, att) = trans(\llbracket I_{EA} \rrbracket(att)(i)) \in J$$

where  $trans$  is defined in Section 4.1, and for every foreign key  $fk : s \rightarrow s' \in S$ ,

$$(i, fk) : i \rightarrow \llbracket I_E \rrbracket(fk)(i) \in \int I \quad F((i, fk)) := fk \quad i.(i, fk) = \llbracket I_E \rrbracket(fk)(i) \in J$$

In addition, for each generator  $g := e.fk_1 \dots fk_n.att$  in  $talg(I)$ , we have a term  $[g] \in Terms(J, \emptyset)$  defined as

$$nf_{I_E}(e) \cdot (nf_{I_E}(e), fk_1) \cdot (nf_{I_E}(e.fk_1), fk_2) \cdot \dots \cdot (nf_{I_E}(e.fk_1, \dots, fk_n), att)$$

and we add  $g = [g]$  to  $J$ . See Figure 18 for an example.

The *co-pivot* of  $I$  is defined to be a schema  $\oint I$  on  $Ty$  that extends  $S$  and an inclusion mapping  $F : S \hookrightarrow \oint I$ . First, we add a single entity  $\star$  to  $\oint I$ , and for every generator  $g : s \in talg(I)$ , an attribute  $g_A$ :

$$\star \in \oint I \quad g_A : \star \rightarrow s \in \oint I$$

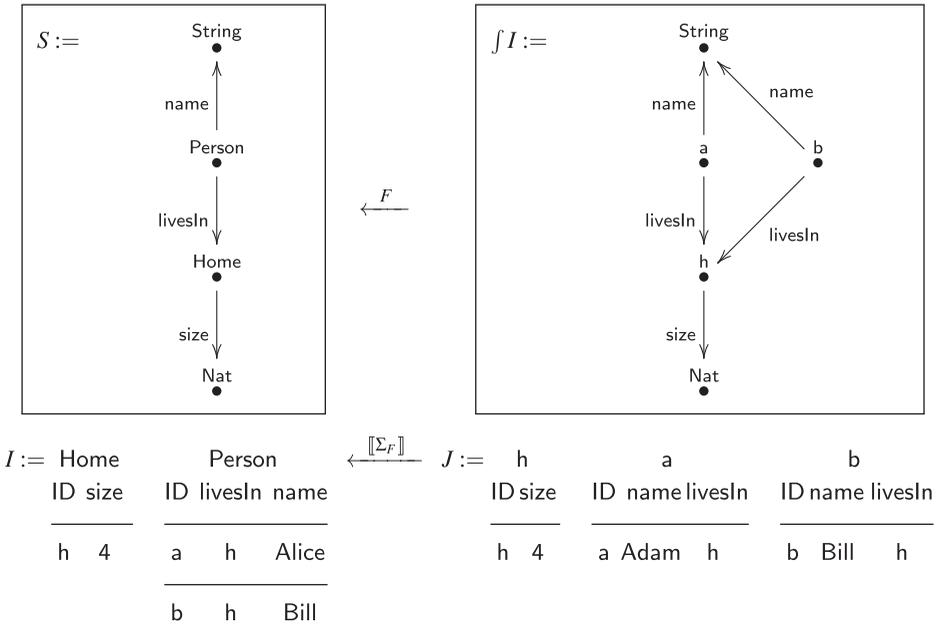


Fig. 18. Example of a pivot.

Then, for every entity  $s \in S$ , and  $i \in \llbracket I_E \rrbracket(s)$ , we add a foreign key:

$$i_E : \star \rightarrow s \in \oint I \quad F(s) := s$$

and additionally, for every attribute  $att : s \rightarrow s'$ , and every foreign key  $fk : s \rightarrow s''$ , we add

$$\forall x : \star. x.i_E.att = x.trans(i.att)_A \quad \forall x : \star. x.i_E.fk = x.nf_{I_E}(i.fk)_E \in \oint I$$

where  $trans$  is defined in Section 4.1. See Figure 19 for an example. Note that in this figure, we include an instance  $J := \Delta_F(I)$  to make the duality of co-pivoting and pivoting explicit, and that although  $J$  includes the tables from  $I$ , we do not display these tables in order to make the figure smaller.

### 5.8 Performance

In this section, we give preliminary performance results for the June 2017 version of the AQL tool. Some of these experiments reference schemas and constructions (e.g., colimits) that are not introduced until the next section (Section 6). In practice, the time required to check schema mappings and uber-flowers for well-formedness (Section 4.3.1) is negligible, so in this section we focus on the scalability of the saturation procedure (Section 5.2; used for  $\Sigma$ ) and uber-flower evaluation (Section 5.4; used for  $\Delta, \Pi$ ).

When interpreting the following results, it is important to consider several caveats. First, the AQL tool is naively implemented; for example, it uses string-valued variables rather than a more sophisticated variable representation such as De

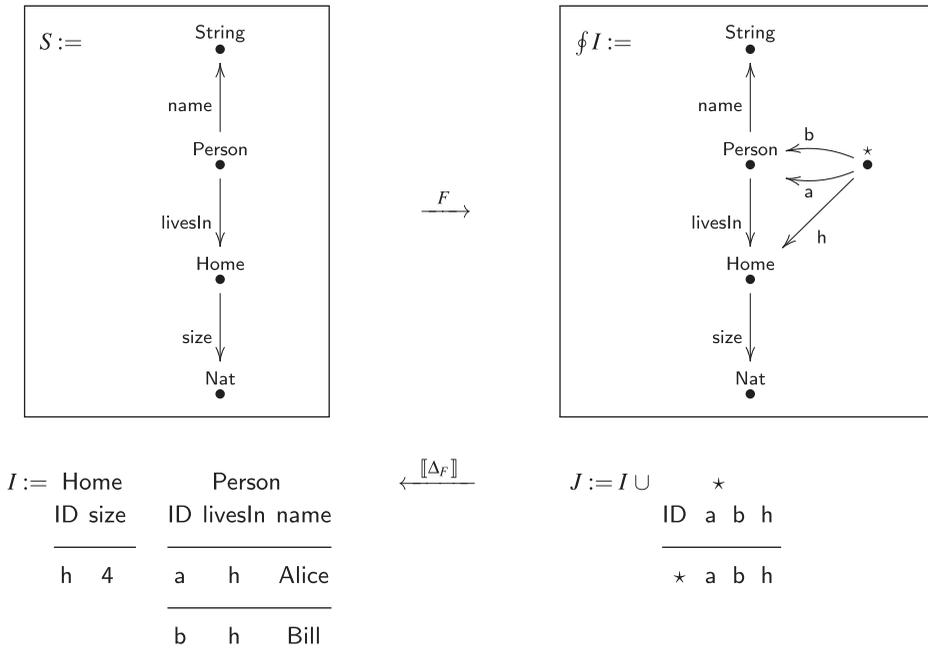


Fig. 19. Example of a co-pivot.

Bruijn indices (Mitchell, 1996). Second, the AQL tool deviates somewhat from the formalism in this paper for reasons of efficiency and expediency. Third, there are other evaluation strategies for AQL besides the one that the AQL tool uses, for example, implementation via SQL generation (Spivak & Wisnesky, 2015) (for  $\Delta$  and  $\Pi$ ) or using the “chase” algorithm (Fagin *et al.*, 2005b) (for  $\Sigma$ ) – in practice, these strategies may be necessary to get practical performance from our algebraic approach to data integration.

As previously mentioned, the AQL tool ships with many different automated theorem provers. In this section, we refer to three of these provers by name:

- *Thue*. This prover implements a Knuth–Bendix completion algorithm specialized to equational theories where all symbols are 0 or 1-ary (Kapur & Narendran, 1985).
- *Congruence*. This prover implements a congruence closure algorithm (Nelson & Oppen, 1980) to decide equational theories without quantifiers.
- *Program*. This prover orients equations into size-reducing rewrite rules and checks for “weak orthogonality” (Baader & Nipkow, 1998). Such theories are first-order functional programs and are decided by rewriting terms into normal forms.

The experiments in this section are synthetic. AQL includes the ability to construct random instances: let  $gens(s)$  indicate a desired set of generators for each entity  $s$  in some schema. For each attribute or foreign key  $f : s \rightarrow s'$  and for each generator  $g \in gens(s)$  the random instance contains an equation  $f(g) = g'$ , where  $g'$  is a (uniformly) randomly chosen element of  $gens(s')$ . These random instances thus have

Total Rows	Thue	Program	Congruence
500	6	1	4
1000	41	10	34
1500	134	32	116
2000	314	78	252
2500	552	156	504
3000	987	189	725
3500	1451	286	1270
4000	2236	446	1773

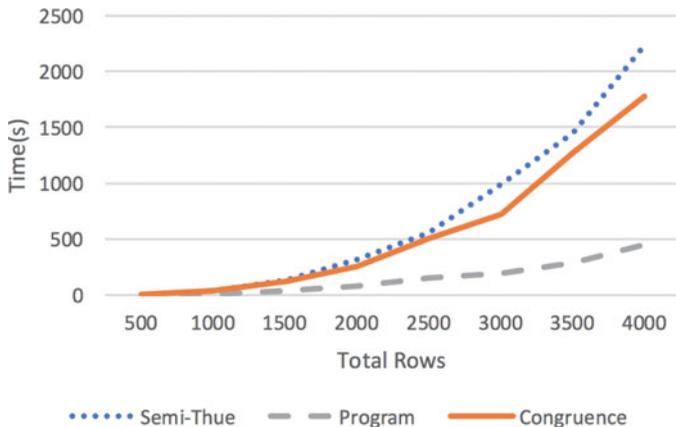


Fig. 20. Time (seconds) to saturate a random instance on the integrated schema of the pharma example (Figure 25), for three theorem provers.

a special, “dense” form, and as such performance tests on them may not indicate real-world performance.

The experiments in this section refer to two kinds of tasks:

- *Saturation*. That is, building a term model/initial algebra from an equational theory by *first constructing a decision procedure for the theory* and then enumerating terms (Section 5.2). Saturation arises following a  $\Sigma$  or pushout operation (see Section 6); data that is imported from, e.g., a SQL database is already saturated. An examination of Java profiling info suggests that the time to construct a decision procedure dominates the time required to build the term model.
- *Query evaluation*. That is, building an instance by evaluating a query on another instance (Section 5.4). Although this task does involve using a decision

Total Rows	Saturation (Thue)	Query Evaluation
90	0.26	0.03
450	9.6	1.2
810	57.5	5.7
1170	263.3	19.6
1530	1150	38

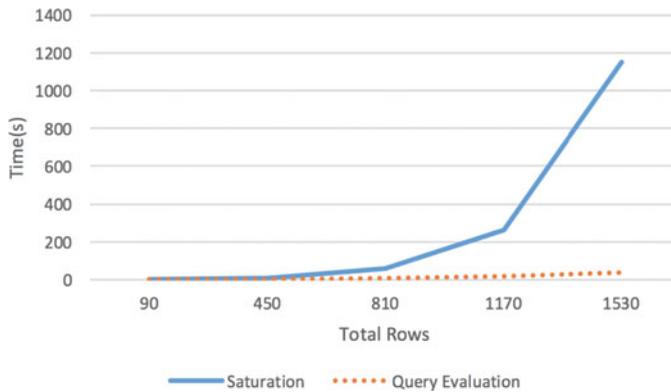


Fig. 21. Time (seconds) to saturate and query a random instance on the colimit schema of the Finance integration example using the thue prover.

procedure, it does not require constructing a new decision procedure from scratch (a decision procedure for the output instance of a  $\Delta$  or a  $\Pi$  or an *eval* can be easily obtained from a decision procedure for the input instance). Hence, performance on this task is tantamount to the performance of the naive nested loops join (Garcia-Molina *et al.*, 2008) algorithm described in Section 5.4. We have also implemented a faster hash join (Garcia-Molina *et al.*, 2008) algorithm, but it is out of scope for this paper.

The results in this section were obtained on an Intel Core i7-6770HQ CPU at 2.60 GHz. All entries are averages: we occasionally saw  $2\times$  speed ups or slow downs from run to run; we are not sure if this is due to AQL’s inherent non-determinism (e.g., the order in which sets are traversed) or some artifact of the Java virtual machine upon which AQL runs.

- *Pharma data integration example.* This example uses the “Pharma Colim” example that is built-in to the AQL IDE and described in Section 6. It shows the amount of time, in seconds, required to saturate a random instance on a given number of rows on the pushout schema in Figure 24. The experiment is repeated for three different provers: Thue, Congruence, and Program. We find that the overhead of the Thue algorithm compared to congruence closure is

Total Rows	Saturation (Thue)	Query Evaluation
300	1.2	0.2
600	5.6	1.1
900	19	3.6
1200	45.3	8.88
1500	84.3	16
1800	147.6	28
2100	256	45.1

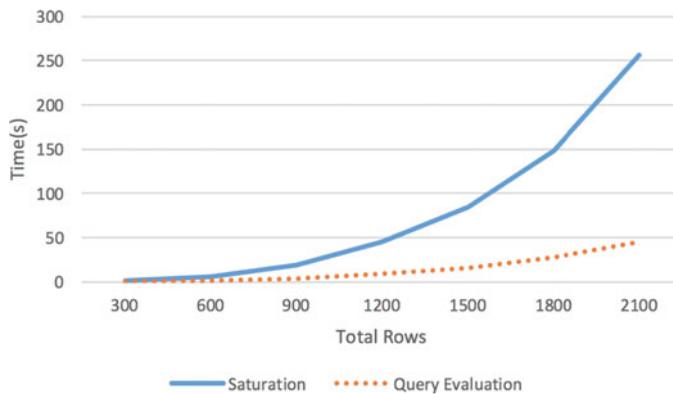


Fig. 22. Time (seconds) to saturate and query a random instance on the pullback schema using the thue prover.

small, and that orienting equations according to size and checking for weak orthogonality is the fastest theorem proving method by a large margin.

- *Finance data integration example.* This example uses the “Finance Colim” example that is built-in to the AQL IDE. It shows the amount of time, in seconds, required to saturate a random instance on a given number of rows on a schema containing 9 entities, 8 foreign keys, 52 attributes, and 12 equations. It also includes an additional schema (9 entities, 16 foreign keys, 50 attributes, 16 equations) and measures the time it takes to evaluate an uber-flower between the two schemas. We find that saturation time dominates query evaluation time.
- *Pullback example.* This example uses the “Pullback” example that is built-in to the AQL IDE. It defines a schema that is a co-span  $\cdot \rightarrow \cdot \leftarrow \cdot$  and a schema that is a commutative square. It saturates a random instance on the co-span schema and then evaluates a query that joins the two root entities into the square schema. We find that saturation time exceeds query evaluation time, but not nearly as much as with the above finance example.

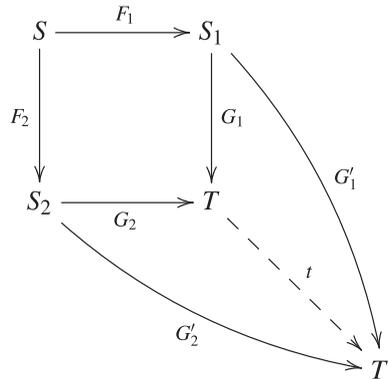


Fig. 23. Pushouts.

## 6 A pushout design pattern for algebraic data integration

In this section, we describe a design pattern for integrating two instances on two different schemas, relative to an overlap schema and an overlap instance, using the formalism defined in this paper. The overlap schema is meant to capture the schema elements common to the two input schemas (e.g., that *Patient* and *Person* should be identified), and the overlap instance is meant to capture the instance data that should be identified (e.g., that *Pete* and *Peter* are the same person).

Although the  $\Sigma, \Delta, \Pi$  data migration functors are sufficient to express queries and data migrations, as unary operations they are insufficient to express data integrations, which involve many schemas and instances and their relationships (Doan *et al.*, 2012). So, we need to define additional operations on our formalism as we develop our pattern. In particular, we will define *pushouts* (Barr & Wells, 1995) of schemas and instances and use pushouts as the basis of our pattern. The idea of using pushouts to integrate data is not new and was for example discussed in Goguen (2004); our goal here is to express this pattern using our formalism. With pushouts defined, we describe our pattern at an abstract level, and then we describe a medical records example that uses the pattern. This example is built into the AQL tool as the “Pharma Colim” example.

Pushouts have a dual, called a pullback, obtained by reversing the arrows in the pushout diagram (Figure 23). Exploring applications of pullbacks to data integration, as well as finding other useful design patterns for algebraic data integration, are important areas for future work.

### 6.1 Pushouts of schemas and instances

Let  $\mathbf{C}$  be a category and  $F_1 : S \rightarrow S_1$  and  $F_2 : S \rightarrow S_2$  be morphisms in  $\mathbf{C}$ . A *pushout* of  $F_1, F_2$  is any pair  $G_1 : S_1 \rightarrow T$  and  $G_2 : S_2 \rightarrow T$  such that  $G_2 \circ F_2 = G_1 \circ F_1$ , with the property that for any other pair  $G'_1 : S_1 \rightarrow T'$  and  $G'_2 : S_2 \rightarrow T'$  for which  $G'_2 \circ F_2 = G'_1 \circ F_1$ , there exists a unique  $t : T \rightarrow T'$  such that  $G'_1 = t \circ G_1$  and  $G'_2 = t \circ G_2$ , as shown in Figure 23.

Our formalism admits pushouts of schemas and instances. Let  $S := (Ens, Symbols, Eqs)$ ,  $S_1 := (Ens_1, Symbols_1, Eqs_1)$  and  $S_2 := (Ens_2, Symbols_2, Eqs_2)$  be schemas on some type-side, where  $Ens$  indicates entities,  $Symbols$  indicates foreign keys and attributes, and  $Eqs$  indicates schema, but not type-side, equations. Let  $F_1 : S \rightarrow S_1$  and  $F_2 : S \rightarrow S_2$  be schema mappings. The pushout schema  $T$  is defined with entities:

$$Ens_T := (Ens_1 \sqcup Ens_2) / \sim$$

where  $\sqcup$  means disjoint union,  $\sim$  is the least equivalence relation such that  $F_1(e) \sim F_2(e)$  for every entity  $e \in S$ , and  $/$  means set-theoretic quotient. We define further that

$$Symbols_T := Symbols_1 \sqcup Symbols_2$$

$$Eqs_T := Eqs_1 \sqcup Eqs_2 \sqcup$$

$$\{v_1 : F_1(s_1), \dots, v_n : F_1(s_n). F_1(e) = F_2(e) : F_1(s) \mid e : s_1 \times \dots \times s_n \rightarrow s \in Symbols_S\}$$

and the schema mappings  $G_1$  and  $G_2$  inject each entity into its equivalence class under  $\sim$  and inject each symbol appropriately.

Pushouts of instances are slightly easier to define than pushouts of schemas. Let  $S := (Gens, Eqs)$ ,  $S_1 := (Gens_1, Eqs_1)$  and  $S_2 := (Gens_2, Eqs_2)$  be instances on some schema, where  $Gens$  indicates generators and  $Eqs$  indicates instance, but not schema, equations. Let  $F_1 : S \rightarrow S_1$  and  $F_2 : S \rightarrow S_2$  be transforms. The pushout instance  $T$  is

$$Gens_T := Gens_1 \sqcup Gens_2$$

$$Eqs_T := Eqs_1 \sqcup Eqs_2 \sqcup \{F_1(e) = F_2(e) : s \mid e : s \in Gens_S\}$$

and the transforms  $G_1, G_2$  are inclusions.

The pushout schemas and instances defined in this section are canonical, but the price for canonicity is that their underlying equational theories tend to be highly redundant (i.e., have many symbols that are provably equal to each other). These canonical pushout schemas and instances can be simplified, and in fact AQL can perform simplification, but the simplification process is necessarily non-canonical (a fact that complicates the use of algebraic specification techniques in general; Mossakowski *et al.*, 2017). In our extended medical example (Figure 24), we will use a simplified non-canonical pushout schema.

### 6.2 Overview of the pattern

Given input schemas  $S_1, S_2$ , an overlap schema  $S$ , and mappings  $F_1, F_2$  as such

$$S_1 \xleftarrow{F_1} S \xrightarrow{F_2} S_2$$

we propose to use their pushout

$$S_1 \xrightarrow{G_1} T \xleftarrow{G_2} S_2$$

as the integrated schema. Given input  $S_1$ -instance  $I_1$ ,  $S_2$ -instance  $I_2$ , overlap  $S$ -instance  $I$  and transforms  $h_1 : \Sigma_{F_1}(I) \rightarrow I_1$  and  $h_2 : \Sigma_{F_2}(I) \rightarrow I_2$ , we propose the

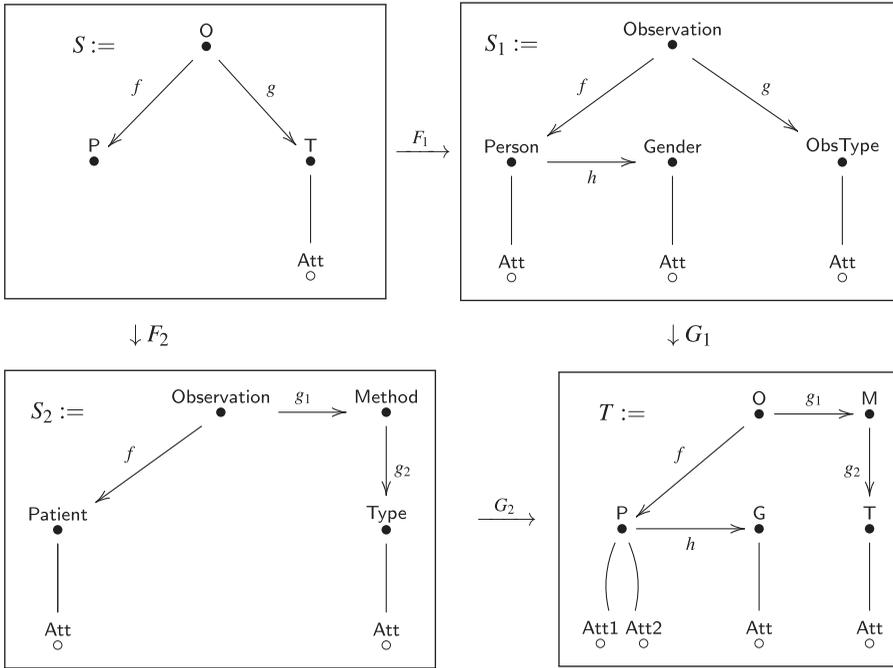


Fig. 24. Medical records schema integration.

pushout of

$$\Sigma_{G_1}(I_1) \xleftarrow{\Sigma_{G_1}(h_1)} (\Sigma_{G_1 \circ F_1}(I) = \Sigma_{G_2 \circ F_2}(I)) \xrightarrow{\Sigma_{G_2}(h_2)} \Sigma_{G_2}(I_2)$$

as the integrated  $T$ -instance.

Because pushouts are initial among the solutions to our design pattern, our integrated instance is the “best possible” solution in the sense that if there is another solution to our pattern, then there will be a unique transform from our solution to the other solution. In functional programming terminology, this means our solution has “no junk” (extra data that should not appear) and “no noise” (missing data that should appear) (Mitchell, 1996). Initial solutions also appear in the theory of relational data integration, where the chase constructs weakly initial solutions to data integration problems (Fagin *et al.*, 2005b).

### 6.3 An example of the pattern

As usual for our formalism, we begin by fixing a type-side. We choose the *Type* type-side from Figure 4. Then, given two source schemas  $S_1, S_2$ , an overlap schema  $S$ , and mappings  $F_1, F_2$  as input, our goal is to construct a pushout schema  $T$  and mappings  $G_1, G_2$ , as shown in Figure 24. In that figure’s graphical notation, an attribute  $\bullet_A \rightarrow_{att} \bullet_{String}$  is rendered as  $\bullet_A - \circ_{att}$ . Next, given input  $S_1$ -instance  $I_1, S_2$ -instance  $I_2$ , overlap  $S$ -instance  $I$  and morphisms  $h_1 : \Sigma_{F_1}(I) \rightarrow I_1$  and  $h_2 : \Sigma_{F_2}(I) \rightarrow I_2$ , our goal is to construct a pushout  $T$ -instance  $J$  and morphisms  $j_1, j_2$ , as shown in Figure 25; note that in this figure, by  $K \rightarrow_{(F,h)} L$  we mean that  $h : \Sigma_F(K) \rightarrow L$ , and

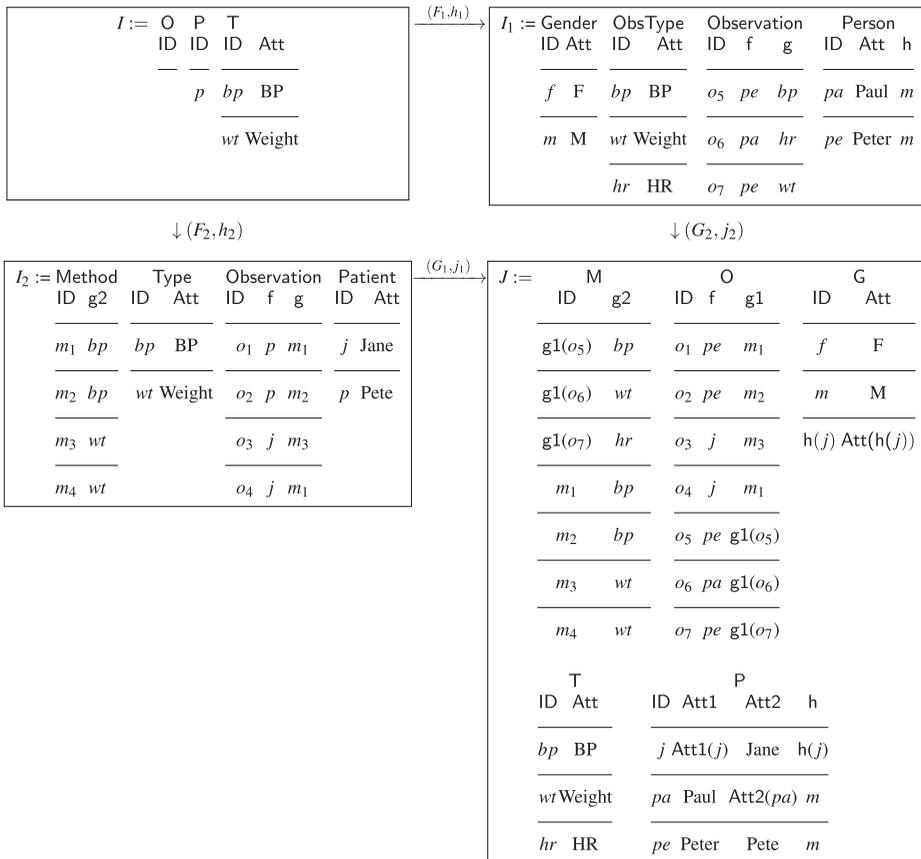


Fig. 25. Medical records data integration.

that we use *italic* font for generators, **sans serif** font for sorts and symbols, and regular font for terms in the type-side.

Our example involves integrating two different patient records databases. In  $S_1$ , the “observations” done on a patient have types, such as heart rate and blood pressure. In  $S_2$ , the observations still have types, but via “methods” (e.g., patient self-report, by a nurse, by a doctor, etc.; for brevity, we have omitted attributes for the names of these methods). Another difference between schemas is that  $S_1$  assigns each patient a gender, but  $S_2$  does not. Finally, entities with the same meaning in both schemas can have different names (Person versus Patient, for example).

We construct the overlap schema  $S$  and mappings  $F_1, F_2$  (Figure 24) by thinking about the meaning of  $S_1$  and  $S_2$ ; alternatively, schema-matching techniques (Doan *et al.*, 2012) can be used to construct overlap schemas. In this example, it is clear that  $S_1$  and  $S_2$  share a common span  $P \xleftarrow{f} O \xrightarrow{g} T$  relating patients, observations, and observation types; in  $S_1$ , this span appears verbatim but in  $S_2$ , the path  $g$  corresponds to  $g_2 \circ g_1$ . This common span defines the action of  $F_1$  and  $F_2$  on entities and foreign keys in  $S$ , so now we must think about the attributes in  $S$ . For purposes of exposition we assume that the names of observation types (“BP,” “Weight,”

etc.) are the same between the instances we are integrating. Hence, we include an attribute for observation type in the overlap schema  $S$ . On the other hand, we do not assume that patients have the same names across the instances we are integrating; for example, we have the same patient named “Pete” in one database and “Peter” in the other database. Hence, we do not include an attribute for patient name in  $S$ . If we did include an attribute for patient name, then the pushout schema would have a single attribute for patient name, and the integrated instance would include the equation “Pete” = “Peter” : String. We would violate the conservative extension property (see Section 5.3), which is not a desirable situation (Ghilardi *et al.*, 2006). So, our design pattern explicitly recommends that when two entities in  $S_1$  and  $S_2$  are identified in an overlap schema, we should only include those attributes that appear in both  $S_1$  and  $S_2$  for which the actual values of these attributes will correspond in the overlap instance. As another example of this phenomenon, to a first approximation, attributes for globally unique identifiers such as social security numbers can be added to overlap schemas, but attributes for non-standard vocabularies such as titles (e.g., CEO versus Chief Executive Officer) should not be added to overlap schemas.

With the overlap schema in hand, we now turn toward our input data. We are given two input instances,  $I_1$  on  $S_1$  and  $I_2$  on  $S_2$ . Entity-resolution (ER) techniques (Doan *et al.*, 2012) can be applied to construct an overlap instance  $I$  automatically. Certain ER techniques can even be implemented as queries in the AQL tool, as we will describe in the next section. However, for the purposes of this example, we will construct the overlap instance by hand. We first assume there are no common observations across the instances; for example, perhaps a cardiologist and nephrologist are merging their records. We also assume that the observation type vocabulary (e.g., “BP” and “Weight”) are standard across the input instances, so we put these observation types into our overlap instance. Finally, we see that there is one patient common to both input instances, and he is named Peter in  $I_1$  and Pete in  $I_2$ , so we add one entry for Pete/Peter in our overlap instance. We have thus completed the input to our design pattern (Figure 25).

In the output of our pattern (Figure 25) we see that the observations from  $I_1$  and  $I_2$  were disjointly unioned together, as desired; that the observation types were (not disjointly) unioned together, as desired; and that Pete and Peter correspond to the same person in the integrated instance. In addition, we see that Jane could not be assigned a gender.

#### 6.4 Pushout-based data integration in practice

In practice in the AQL tool, we observe several phenomena that are not accounted for in the above theoretical description of the pushout pattern:

- In those data integration scenarios where a desired integrated schema is given (e.g., a particular star schema motivated by analytics), the process of migrating data from the pushout schema to the target schema seems to invariably be

given by a  $\Pi \circ \Delta$  migration; i.e., be given by evaluating (rather than co-evaluating) uber-flowers.

- The integrated (pushed-out) instance may more “baggy” than desired; i.e., there may be multiple rows that are equivalent up to the values of their attributes and foreign keys. The de-duplication operation discussed in Spivak & Wisnesky (2015) can be used to “make distinct” the rows of the integrated instance.
- When generalizing from pushouts to arbitrary colimits (essentially,  $n$ -ary pushouts; Barr & Wells, 1995) it becomes apparent that there is significant redundancy in the above span-of-mappings approach to specifying colimits at both the schema and instance levels. We can avoid this redundancy by specifying colimits by means of co-products (disjoint unions) and quotients. For example, AQL pseudo-code for a quotient-based specification of the schema pushout in Figure 24 is

```

schema T = S1 + S2 /
  entity equations
    S1_Observation = S2_Observation
    S1_Person = S2_Patient
    S1_ObsType = S2_Type
  path equations
    S1_f = S2_f
    S1_g = S2_g1.S2_g2
    S1_ObsType_att = S2_Type_att

```

The above AQL code fragment is six times shorter than the corresponding AQL code for specifying pushouts using spans of schema mappings, and we now always prefer the quotient based approach in practice. Quotients can also be used to reduce verbosity for instance pushouts.

- As mentioned earlier, finding a “good” (rather than canonical) presentation of a pushout schema is a subtly difficult problem that occurs in similar form in many places in algebraic specification (Mossakowski *et al.*, 2017). As an alternative to pushouts, so-called pseudo-pushouts may be used. Conceptually, these are similar to pushouts, but entities are *made isomorphic* rather than *equated*. Although pseudo-pushout schemas are “maximally large” (the number of entities in the pseudo pushout of  $S_1$  and  $S_2$  will be the sum of the number of entities in  $S_1$  and  $S_2$ ), the entities, attributes, and foreign keys of pseudo-pushout schemas are very natural to name. Comparing pushouts to pseudo-pushouts for data integration purposes is an important area for future work.

### 6.5 Entity-resolution using uber-flowers

Let schemas  $S, S_1, S_2$ , mappings  $F_1 : S \rightarrow S_1, F_2 : S \rightarrow S_2$ , and  $S_1$ -instance  $I_1$  and  $S_2$ -instance  $I_2$  be given. In practice, we anticipate that sophisticated *entity-resolution* (Doan *et al.*, 2012) techniques will be used to construct the overlap  $S$ -instance  $I$  and transforms  $h_1 : \Sigma_{F_1}(I) \rightarrow I_1$  and  $h_2 : \Sigma_{F_2}(I) \rightarrow I_2$ . However, it

is possible to perform a particularly simple kind of entity resolution directly by evaluating uber-flowers.

Technically, the overlap instance used in the pushout pattern should not be thought of as containing resolved (unified) entities; rather, it should be thought of as containing the *record linkages* between entities that will resolve (unify) (Doan *et al.*, 2012). The pushout resolves entities by forming equivalence classes of entities under the equivalence relation induced by the links. As the size of the overlap instance gets larger, the size of the pushout gets smaller, which is the opposite of what would happen if the overlap instance contained the resolved entities themselves, rather than the links between them. For example, let  $A$  and  $B$  be instances on some schema that contains an entity `Person`, and let  $A(\text{Person}) := \{a_1, a_2\}$  and  $B(\text{Person}) := \{b_1, b_2\}$ . If the overlap instance  $O$  has  $O(\text{Person}) := \{(a_1, b_1), (a_1, b_2), (a_2, b_1), (a_2, b_2)\}$ , then this does not mean that the pushout will have four people; rather, the pushout will have one person corresponding to  $\{a_1, a_2, b_1, b_2\}$ , because these four people are linked. Intuitively, the overlap instance  $I$  constructed by the technique in this section is (isomorphic to) a sub-instance of  $\Delta_{F_1}(I_1) \times \Delta_{F_2}(I_2)$ , where  $\times$  denotes a kind of product of instances that we will not define here.

Let  $inc_1 : S_1 \rightarrow S_1 + S_2$  and  $inc_2 : S_2 \rightarrow S_1 + S_2$  be inclusion schema mappings, and define the  $S_1 + S_2$  instance  $I' := \Sigma_{inc_1}(I_1) + \Sigma_{inc_2}(I_2)$ . This instance will contain  $I_1$  and  $I_2$  within it, and will contain nothing else. (Here,  $X + Y$  means co-product, which is equivalent to the pushout of  $X$  and  $Y$  over the empty schema or instance.)

We construct overlap  $S$ -instance  $I$  by defining a query  $Q : S_1 + S_2 \rightarrow S$  and evaluating it on the  $S_1 + S_2$ -instance  $I'$ . For each entity  $s \in S$ , we choose a set of pairs of attributes from  $F_1(s)$  and  $F_2(s)$  that we desire to be “close.” In the medical records example, for  $P$  we choose (`PatientAtt`, `PersonAtt`) and for  $T$  we choose (`ObsTypeAtt`, `TypeAtt`); we choose nothing for  $O$ . We next choose a way to compare these attributes; for example, we choose a string edit distance of less than two to indicate that the entities match. This comparison function must be added to our type-side, e.g.,

$$\text{strMatches} : \text{String} \times \text{String} \rightarrow \text{Nat} \quad \text{true} : \text{Nat} \quad \text{true} = 1$$

The function `strComp` can be defined using equations, although the AQL tool allows such functions to be defined using Java code (see Section 5). With the `String-comparator` in hand, we can now define  $Q : S_1 + S_2 \rightarrow S$  as in Figure 26. The overlap instance  $I$  is defined as  $eval(Q)(I')$ . To construct  $h_n : \Sigma_{F_n}(I) \rightarrow I_n$  for  $n = 1, 2$ , we define projection queries  $Q_n : S_1 + S_2 \rightarrow S$  and inclusion query morphisms  $q_n : Q_n \rightarrow Q$  as in Figure 26 as follows. We start with the induced transforms for  $q_n$ , then apply  $\Sigma_{F_n}$ , then compose with the isomorphism  $eval(Q_n)(I') \cong \Delta_{F_n}(I_n)$ , and then compose the co-unit  $\epsilon$  of the  $\Sigma_{F_n} \dashv \Delta_{F_n}$  adjunction, to obtain  $h_n$  as in Figure 26.

The result of running Figure 26 on the medical records data  $I_1, I_2$  from Figure 25 is the overlap instance  $I$  from Figure 25. To compute the isomorphism  $eval(Q_n)(I') \rightarrow \Delta_{F_n}(I_n)$ , we note that the generators of  $eval(Q_n)(I')$  will be singleton substitutions such as  $[v_n \mapsto inj_n a_n]$  where  $a_n$  is a term in  $\Sigma_{inc_n}(I_n)$  and  $inj_n$  means co-product injection. But  $inc_n$  is an inclusion, so  $a_n$  is a term in  $I_1$ . Because we compute  $\Delta_{F_n}$  by translation into an uber-flower similar to  $Q_n$ , the generators of  $\Delta_{F_n}(I')$  will have a

```

Q1 : S1 + S2 → S :=
  O := for o1 in Observation1
        keys f->[p1->o1.f], g->[t1->o1.g]
  P := for p1 in Person
  T := for t1 in ObsType, return att->t1.att
q1 ↓

Q : S1 + S2 → S :=
  O := for o1 in Observation1, o2 in Observation2
        where 1 = 2
        keys f->[p1->o1.f, p2 -> o2.f],
              g->[t1->o1.g, t2 -> o2.g1.g2]
  P := for p1 in Person, p2 in Patient
        where true = strMatches(p1.PatientAtt, p2.PersonAtt)
  T := for t1 in ObsType, t2 in Type
        where t1.ObsTypeAtt = t2.TypeAtt
        return att->t1.ObsTypeAtt
q2 ↑

Q2 : S1 + S2 → S :=
  O := for o2 in Observation2
        keys f->[p1->o2.f], g->[t1->o2.g1.g2]
  P := for p2 in Patient
  T := for t2 in Type, return att->t2.att
    
```

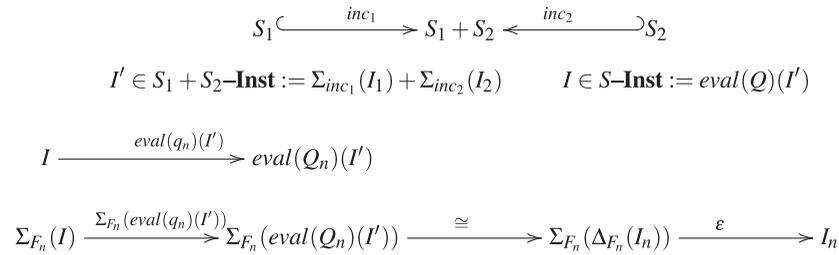


Fig. 26. Entity-resolution on medical records using uber-flowers.

similar form:  $[v'_n \mapsto a_n]$  that defines the necessary isomorphism. When all schemas are disjoint and variables are chosen appropriately, the isomorphism can be made an equality.

### 7 Conclusion

In this paper, we have described an algebraic formalism for integrating data, and work continues. In the short term, we aim to formalize our experimental “computational type-sides,” and to develop a better conservativity checker. In the long term, we are looking to develop other design patterns for data integration and to study their compositions, and we are developing an equational theorem prover tailored to our needs. In addition to these concrete goals, we believe there is much to be gained from the careful study of the differences between our formalism, with its category-theoretic semantics, and the formalism of EDs, with its relational semantics (Doan *et al.*, 2012). For example, there is a semantic similarity between our  $\Sigma$  operation and the chase; as another example, so far we have found no relational counterpart to the concept of query “co-evaluation;” and finally, our uber-flower queries may suggest generalizations of comprehension syntax (Grust, 2004).

## References

- Abiteboul, S., Hull, R. & Vianu, V. (1995) *Foundations of Databases*. Addison-Wesley-Longman.
- Adámek, J., Rosický, J. & Vitale, E. M. (2011) *Algebraic Theories*. Cambridge Tracts in Mathematics, vol. 184. Cambridge University Press.
- Alagic, S. & Bernstein, P. (2001) A model theory for generic schema management. In Proceedings of the 8th International Workshop on Database Programming Languages, 228–246.
- Angles, R. & Gutierrez, C. (2008) Survey of graph database models. *ACM Comput. Surv.*, **40**(1).
- Baader, F. & Nipkow, T. (1998) *Term Rewriting and All That*. Cambridge University Press.
- Bachmair, L., Dershowitz, N. & Plaisted, D. A. (1989) Completion without failure. *Resolution Equ. Algebr. Struct. – Rewriting Tech.* **2**.
- Barr, M. & Wells, C. (1995) *Category Theory for Computing Science*. Prentice Hall International.
- Bertot, Y. & Castéran, P. (2010) *Interactive Theorem Proving and Program Development: Coq'art the Calculus of Inductive Constructions*. Springer.
- Blum, E. K., Ehrig, H. & Parisi-Presicce, F. (1987) Algebraic specification of modules and their basic interconnections. *J. Comput. Syst. Sci.* **34**(2-3), 293–339.
- Bush, M. R., Leeming, M. & Walters, R. F. C. (2003) Computing left Kan extensions. *J. Symbol. Comput.* **35**(2), 107–126.
- Doan, H., Halevy, A. & Ives, Z. (2012) *Principles of Data Integration*. Morgan Kaufmann.
- Enderton, H. B. (2001) *A Mathematical Introduction to Logic*. Academic Press.
- Fagin, R., Kolaitis, P. G., Miller, R. J. & Popa, L. (2005b) Data exchange: Semantics and query answering. *Theoretical Computer Science* **336**(1), 89–124.
- Fagin, R., Kolaitis, P. G., Popa, L. & Tan, W. (2005a) Composing schema mappings: Second-order dependencies to the rescue. *ACM Transactions on Database Systems* **30**(4), 994–1055.
- Fleming, M., Gunther, R. & Rosebrugh, R. (2003) A database of categories. *J. Symbol. Comput.* **35**(2), 127–135.
- Garcia-Molina, H., Ullman, J. D. & Widom, J. (2008) *Database Systems: The Complete Book*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall.
- Ghilardi, S., Lutz, C. & Wolter, F. (2006) Did I damage my ontology? A case for conservative extensions in description logics. In Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning, 187–197.
- Goguen, J. A. & Burstall, R. M. (1984) Introducing institutions. In Proceedings of the Carnegie Mellon Workshop on Logic of Programs, 221–256.
- Goguen, J. (2004) Information Integration in Institutions [online]. Available at: <http://cseweb.ucsd.edu/~goguen/pps/if04.pdf>
- Grust, T. (2004) Monad comprehensions: A versatile representation for queries. In *The Functional Approach to Data Management: Modeling, Analyzing and Integrating Heterogeneous Data*, Gray, P. M. D., Kerschberg, L., King, P. J. H. & Poulouvassilis, A. (eds). Berlin, Heidelberg: Springer, 288–311.
- Haas, L. M., Hernández, M. A., Ho, H., Popa, L. & Roth, M. (2005) Clio grows up: From research prototype to industrial tool. In Proceedings of the International Conference on Management of Data, 805–810.
- Johnson, M., Rosebrugh, R. & Wood, R. J. (2002) Entity-Relationship-Attribute designs and sketches. *Theory Appl. Categories* **10**, 94–112.

- Kapur, D. & Narendran, P. (1985) The Knuth–Bendix completion procedure and Thue systems. *SIAM J. Comput.* **14**(4), 1052–1072.
- Knuth, D. & Bendix, P. (1970) Simple word problems in universal algebra. In *Computational Problems in Abstract Algebra*, Leech, J. (ed). Springer.
- Lüth, C., Roggenbach, M. & Schröder, L. (2005) CCC – The CASL consistency checker. In *Recent Trends in Algebraic Development Techniques, 17th International Workshop*, Fiadeiro, J. (ed). Lecture Notes in Computer Science, vol. 3423, Springer, 94–105.
- Melnik, S., Rahm, E. & Bernstein, P. A. (2003) Rondo: A programming platform for generic model management. In *Proceedings of the International Conference on Management of Data*, 193–204.
- Mitchell, J. C. (1996) *Foundations of Programming Languages*. MIT Press.
- Mossakowski, T., Krumnack, U. & Maibaum, T. (2014) What is a derived signature morphism? In *Recent Trends in Algebraic Development Techniques*, Lecture Notes in Computer Science, vol 9463, Springer, 90–109.
- Mossakowski, T., Rabe, F. & Mihai, C. (2017) Canonical selection of colimits, Algebraic Model Management: A Survey. arXiv:1705.09363.
- Nelson, G. & Oppen, D. C. (1980) Fast decision procedures based on congruence closure. *J. ACM* **27**(2), 356–364.
- Patterson, E. (2017) Knowledge representation in bicategories of relations. arXiv:1706.00526.
- Roberson, E. L. & Wyss, C. M. (2004) Optimal tuple merge is NP-complete. Bloomington: Indiana University School of Informatics and Computing.
- Schultz, P., Spivak, D. I., Vasilakopoulou, C. & Wisnesky, R. (2017) Algebraic databases. *Theory Appl. Categ.* **32**.
- Schultz, P., Spivak, D. I. & Wisnesky, R. (2015) QINL: Query-integrated languages. arXiv: 1511.06459.
- Schultz, P., Spivak, D. I. & Wisnesky, R. (2016) Algebraic model management: A survey. In *Proceedings of the Workshop on Algebraic Development Techniques*.
- Shipman, D. W. (1981) The functional data model and the data languages dplex. *ACM Trans. database Syst.* **6**(1), 140–173.
- Spivak, D. I. (2012) Functorial data migration. *Inform. Comput.* **217**, 31–51.
- Spivak, D. I. (2014) Database queries and constraints via lifting problems. *Math. Struct. Comput. Sci.* **24**(6).
- Spivak, D. I. & Wisnesky, R. (2015) Relational foundations for functorial data migration. In *Proceedings of the 15th Symposium on Database Programming Languages*, 21–28.
- Zieliński, B., Maślanka, P. & Sobieski, Ś. (2013) Allegories for database modeling. In *Model and Data Engineering: 3rd International Conference*, 278–289.