

Implementing Mapping Composition

Philip A. Bernstein
Microsoft Research
philbe@microsoft.com

Todd J. Green*
U. of Pennsylvania
tjgreen@cis.upenn.edu

Sergey Melnik
Microsoft Research
melnik@microsoft.com

Alan Nash*
U.C. San Diego
anash@cs.ucsd.edu

ABSTRACT

Mapping composition is a fundamental operation in metadata driven applications. Given a mapping over schemas σ_1 and σ_2 and a mapping over schemas σ_2 and σ_3 , the composition problem is to compute an equivalent mapping over σ_1 and σ_3 . We describe a new composition algorithm that targets practical applications. It incorporates view unfolding. It eliminates as many σ_2 symbols as possible, even if not all can be eliminated. It covers constraints expressed using arbitrary monotone relational operators and, to a lesser extent, non-monotone operators. And it introduces the new technique of left composition. We describe our implementation, explain how to extend it to support user-defined operators, and present experimental results which validate its effectiveness.

1. INTRODUCTION

A mapping is a relationship between the instances of two schemas. Some common types of mappings are relational queries, relational view definitions, global-and-local-as-view (GLAV) assertions, XQuery queries, and XSL transformations. The manipulation of mappings is at the core of many important data management problems, such as data integration, database design, and schema evolution. Hence, general-purpose algorithms for manipulating mappings have broad application to data management.

Data management problems like those above often require that mappings be composed. The composition of a mapping m_{12} between schemas σ_1 and σ_2 and a mapping m_{23} between schemas σ_2 and σ_3 is a mapping between σ_1 and σ_3 that captures the same relationship between σ_1 and σ_3 as m_{12} and m_{23} taken together.

Given that mapping composition is useful for a variety of database problems, it is desirable to develop a general-purpose composition component that can be reused in many application settings, as was proposed in [1, 3]. This paper reports on the development of such a component, an implementation of a new algorithm for composing mappings between relational schemas. Compared to past approaches, the algorithm handles more expressive mappings, makes a best-effort when it cannot obtain a perfect answer, includes several new heuristics, and is designed to be extensible.

*Work performed during an internship at Microsoft Research.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permission from the publisher, ACM.

VLDB '06, September 12-15, 2006, Seoul, Korea.

Copyright 2006 VLDB Endowment, ACM 1-59593-385-9/06/09.

1.1 Applications of Mapping Composition

Composition arises in many practical settings. In data integration, a query needs to be composed with a view definition. If the view definition is expressed using global-as-view (GAV), then this is an example of composing two functional mappings: a view definition that maps a database to a view, and a query that maps a view to a query result. The standard approach is view unfolding, where references to the view in the query are replaced by the view definition [10]. View unfolding is simply function composition, where a function definition (i.e., the body of the view) is substituted for references to the function (i.e., the view schema) in the query.

In peer-to-peer data management, composition is used to support queries and updates on peer databases. When two peer databases are connected through a sequence of mappings between intermediate peers, these mappings can be composed to relate the peer databases directly. In Piazza [11], such composed mappings are used to reformulate XML queries. In the ORCHESTRA collaborative data sharing system [12], updates are propagated using composed mappings to avoid materializing intermediate relations.

A third example is schema evolution, where a schema σ_1 evolves to become a schema σ'_1 . The relationship between σ'_1 and σ_1 can be described by a mapping. After σ_1 has evolved, any existing mappings involving σ_1 , such as a mapping from σ_1 to schema σ_2 , must now be upgraded to a mapping from σ'_1 to σ_2 . This can be done by composing the σ'_1 - σ_1 mapping with the σ_1 - σ_2 mapping. Depending on the application, one or both of these mappings may be non-functional, in which case composing mappings is no longer simply function composition.

A different schema evolution problem arises when an initial schema σ_1 is modified by two independent designers, producing schemas σ_2 and σ_3 . To merge them into a single schema, we need a mapping between σ_2 and σ_3 that describes their overlapping content [4, 9]. This σ_2 - σ_3 mapping can be obtained by composing the σ_1 - σ_2 and σ_1 - σ_3 mappings. Even if the latter two mappings are functions, one of them needs to be inverted before they can be composed. Since the inverse of a function may not be a function, this too entails the composition of non-functional mappings.

Finally, consider a database design process that evolves a schema σ_1 via a sequence of incremental modifications. This produces a sequence of mappings between successive versions of the schema, from σ_1 to σ_2 , then to σ_3 , and so forth, until the desired schema σ_n is reached. At the end of this process, a mapping from σ_1 to the evolved schema σ_n is needed, for example, as input to the schema evolution scenarios above. This mapping can be obtained by composing the mappings between the successive versions of the schema. The following example illustrates this last scenario.

EXAMPLE 1 Consider a schema editor, in which the designer modifies a database schema, resulting in a sequence of schemas with

mappings between them. She starts with the schema

Movies(mid, name, year, rating, genre, theater)

where *mid* means movie identifier. The designer decides that only 5-star movies and no ‘theater’ or ‘genre’ should be present in the database; she edits the table obtaining the following schema and mapping:

FiveStarMovies(mid, name, year)

$$\pi_{\text{mid,name,year}}(\sigma_{\text{rating}=5}(\mathbf{Movies})) \subseteq \mathbf{FiveStarMovies} \quad (1)$$

To improve the organization of the data, the designer then splits the **FiveStarMovies** table into two tables, resulting in a new schema and mapping

Names(mid, name) **Years**(mid, year)

$$\pi_{\text{mid,name,year}}(\mathbf{FiveStarMovies}) \subseteq \mathbf{Names} \bowtie \mathbf{Years} \quad (2)$$

The system composes mappings (1) and (2) into a new mapping:

$$\pi_{\text{mid,name}}(\sigma_{\text{rating}=5}(\mathbf{Movies})) \subseteq \mathbf{Names}$$

$$\pi_{\text{mid,year}}(\sigma_{\text{rating}=5}(\mathbf{Movies})) \subseteq \mathbf{Years}$$

With this mapping, the designer can now migrate data from the old schema to the new schema, reformulate queries posed over one schema to equivalent queries over the other schema, etc. \square

1.2 Related Work

Mapping composition is a challenging problem. Madhavan and Halevy [6] showed that the composition of two given mappings expressed as GLAV formulas may not be expressible in a finite set of first-order constraints. Fagin, Kolaitis, Popa, and Tan [5] showed that the composition of certain kinds of first-order mappings may not be expressible in any first-order language, even by an infinite set of constraints. That is, that language is not *closed* under composition. Nash, Bernstein, and Melnik [8] showed that for certain classes of first-order languages, it is undecidable to determine whether there is a finite set of constraints in the same language that represents the composition of two given mappings. These results are sensitive to the particular class of mappings under consideration. But in all cases the mapping languages are first-order and are therefore of practical interest.

In [5] Fagin et al. introduced a second-order mapping language that is closed under composition, namely second-order source-to-target tuple-generating dependencies. A second-order language is one that can quantify over function and relation symbols. A tuple-generating dependency specifies an inclusion of two conjunctive queries, $Q_1 \subseteq Q_2$. It is called source-to-target when Q_1 refers only to symbols from the source schema and Q_2 refers only to symbols from the target schema. The second-order language of [5] uses existentially quantified function symbols, which essentially can be thought of as Skolem functions. Fagin et al. present a composition algorithm for this language and show it can have practical value for some data management problems, such as data exchange. However, using it for that purpose requires a custom implementation, since the language of second-order tuple-generating dependencies is not supported by standard SQL-based database tools.

Yu and Popa [14] considered mapping composition for second-order source-to-target constraints over nested relational schemas in support of schema evolution. They presented a composition algorithm similar to the one in [5], with extensions to handle nesting, and with significant attention to minimizing the size of the result. They reported on a set of experiments using mappings on both synthetic and real life schemas, to demonstrate that their algorithm is fast and is effective at minimizing the size of the result.

Nash et al. [8] studied the composition of first-order constraints that are not necessarily source-to-target. They consider dependencies that can express key constraints and inclusions of conjunctive queries $Q_1 \subseteq Q_2$ where Q_1 and Q_2 may reference symbols from both the source and target schema. They do not allow existential quantifiers over function symbols. The composition of constraints in this language is not closed and determining whether a composition result exists is undecidable. Nevertheless, they gave an algorithm that produces a composition, if it halts (which it may not do).

Like Nash et al. [8], we explore the mapping composition problem for constraints that are not restricted to being source-to-target. Our algorithm strictly extends that of Nash et al. [8], which in turn strictly extends that of Fagin et al. [5] for source-to-target embedded dependencies. If the input is a set of source-to-target embedded dependencies, our algorithm behaves similarly to that in [5], except that as in [8] we also attempt to express the result as embedded dependencies through a deskolemization step. It is known from results in [5] that such a step can not always succeed. Furthermore, we also apply a ‘‘left-compose’’ step which allows the algorithm to handle mappings on which the algorithm in [8] fails.

1.3 Contributions

Given the inherent difficulty of the problem and limitations of past approaches, we recognized that compromises and special features would be needed to produce a mapping composition algorithm of practical value. The first issue was which language to choose.

Algebra-based rather than logic-based. We wanted our composition algorithm to be directly usable by existing database tools. We therefore chose a relational algebraic language: Each mapping is expressed as a set of constraints, each of which is either a containment or equality of two relational algebraic expressions. This language extends the algebraic dependencies of [13]. Each constraint is of the form $E_1 = E_2$ or $E_1 \subseteq E_2$ where E_1 and E_2 are arbitrary relational expressions containing not only select, project, and join but possibly many other operators. Calculus-based languages have been used in all past work on mapping composition we know of. We chose relational algebra because it is the language implemented in all relational database systems and most tools. It is therefore familiar to the developers of such systems, who are the intended users of our component. It also makes it easy to extend our language simply by allowing the addition of new operators. Notice that the relational operators we handle are sufficient to express embedded dependencies. Therefore, the class of mappings which our algorithm accepts includes embedded dependencies and, by allowing additional operators such as set difference, goes beyond them.

Eliminates one symbol at a time. Our algorithm for composing these types of algebraic mappings gives a partial solution when it is unable to find a complete one. The heart of our algorithm is a procedure to eliminate relation symbols from the intermediate signature σ_2 . Such elimination can be done one symbol at a time. Our algorithm makes a best effort to eliminate as many relation symbols from σ_2 as possible, even if it cannot eliminate all of them. By contrast, if the algorithm in [8] is unable to produce a mapping over σ_1 and σ_3 with no σ_2 -symbols, it simply runs forever or gives up. In some cases it may be better to eliminate some symbols from σ_2 successfully, rather than insist on either eliminating all of them or failing. Thus, the resulting mapping may be over σ_1 , σ'_2 , and σ_3 , where σ'_2 is a subset of σ_2 instead of over just σ_1 and σ_3 .

To see the value of this best-effort approach, consider a composition that produces a mapping m that contains a σ_2 -symbol S . If m is later composed with another mapping, it is possible that the latter composition can eliminate S . (We will see examples of this

later, in our experiments.) Also, the inability to eliminate S may be inherent in the given mappings. For example, S may be involved in a recursive computation that cannot be expressed purely in terms of σ_1 and σ_3 such those of Theorem 1 in [8]:

$$R \subseteq S, \quad S = \text{tc}(S), \quad S \subseteq T$$

where $\sigma_1 = \{R\}$, $\sigma_2 = \{S\}$, $\sigma_3 = \{T\}$ with R, S, T binary and where the middle constraint says that S is transitively closed. In this case, S cannot be eliminated, but is definable as a recursive view on R and can be added to σ_1 . To use the mapping, those non-eliminated σ_2 -symbols may need to be populated as intermediate relations that will be discarded at the end. In this example this involves low computational cost. In many applications it is better to have such an approximation to a desired composition mapping than no mapping at all. Moreover, in many cases the extra cost associated with maintaining the extra σ_2 symbols is low.

Tolerance for unknown or partially known operators. Instead of rejecting an algebraic expression because it contains unknown operators which we do not know how to handle, our algorithm simply delays handling such operators as long as possible. Sometimes, it needs no knowledge at all of the operators involved. This is the case, for example, when a subexpression that contains an unknown operator can be replaced by another expression. At other times, we need only partial knowledge about an operator. Even if we do not have the partial knowledge we need, our algorithm does not fail globally, but simply fails to eliminate one or more symbols that perhaps it could have eliminated if it had additional knowledge about the behavior of the operator.

Use of monotonicity. One type of partial knowledge that we exploit is monotonicity of operators. An operator is monotone in one of its relation symbol arguments if, when tuples are added to that relation, no tuples disappear from its output. For example, select, project, join, union, and semijoin are all monotone. Set difference (e.g., $R - S$) and left outerjoin are monotone in their first argument (R) but not in their second (S). Our key observation is that when an operator is monotone in an argument, that argument can sometimes be replaced by an expression from another constraint. For example, if we have $E_1 \subseteq R$ and $E_2 \subseteq S$, then in some cases it is valid to replace R by E_1 in $R - S$, but not to replace S by E_2 . Although existing algorithms only work with select project, join and union, this observation enables our algorithm to handle outerjoin, set difference, and anti-semijoin. Moreover, our algorithm can handle nulls and bag semantics in many cases.

Normalization and denormalization. We call left-normalization the process of bringing the constraints to a form where a relation symbol S that we are trying to eliminate appears in a single constraint alone on the left. The result is of the form $S \subseteq E$ where E is an expression. We define right-normalization similarly. Normalization may introduce “pseudo-operators” such as Skolem functions which then need to be eliminated by a denormalization step. Currently we do not do much left-normalization. Our right-normalization is more sophisticated and, in particular, can handle projections by Skolemization. The corresponding denormalization is very complex. An important observation here is that normalization and denormalization are general steps which may possibly be extended on an operator-by-operator basis.

Left compose. One way to eliminate a relation symbol S is to replace S 's occurrences in some constraints by the expression on the other side of a constraint that is normalized for S . There are two versions of this replacement, right compose and left compose. In right compose, we use a constraint $E \subseteq S$ that is right-normalized for S and substitute E for S on the left side of a constraint that is monotonic in S , such as transforming $R \times S \subseteq T$ into $R \times E \subseteq T$,

thereby eliminating S from the constraint. Right composition is an extension of the algorithms in [5, 8]. We also introduce left compose, which handles some additional cases where right compose fails. Suppose we have the constraints $E_2 \subseteq M(S)$ and $S \subseteq E_1$, where $M(S)$ is an expression that is monotonic in S but which we either do not know how to right-normalize or which would fail to right-denormalize. Then left compose immediately yields $E_2 \subseteq M(E_1)$.

Extensibility and modularity. Our algorithm is extensible by allowing additional information to be added separately for each operator in the form of information about monotonicity and rules for normalization and denormalization. Many of the steps are rule-based and implemented in such a way that it is easy to add rules or new operators. Therefore, our algorithm can be easily adapted to handle additional operators without specialized knowledge about its overall design. Instead, all that is needed is to add new rules.

Experimental study. We implemented our algorithm and ran experiments to study its behavior. We used composition problems drawn from the recent literature [5, 7, 8], and a set of large synthetic composition tasks, in which the mappings were generated by composing sequences of elementary schema modifications. We used these mappings to test the scalability and effectiveness of our algorithm in a systematic fashion. Across a range of composition tasks, it eliminated 50-100% of the symbols and usually ran in under a second. We see this study as a step toward developing a benchmark for composition algorithms.

The rest of the paper is organized as follows. Section 2 presents notation needed to describe the algorithm. Section 3 describes the algorithm itself, starting with a high-level description and then drilling into the details of each step, one-by-one. Section 4 presents our experimental results. Section 5 is the conclusion.

2. PRELIMINARIES

We adopt the unnamed perspective which references the attributes of a relation by index rather than by name. A *relational expression* is an expression formed using base relations and the *basic operators* union \cup , intersection \cap , cross product \times , set difference $-$, projection π and selection σ as follows. The name S of a relation is a relational expression. If E_1 and E_2 are relational expressions, then so are

$$\begin{array}{lll} E_1 \cup E_2 & E_1 \cap E_2 & E_1 \times E_2 \\ E_1 - E_2 & \sigma_c(E_1) & \pi_I(E_1) \end{array}$$

where c is an arbitrary boolean formula on attributes (identified by index) and constants and I is a list of indexes. The meaning of a relational expression is given by the standard set semantics. To simplify the presentation in this paper, we focus on these basic six relational operators and view the join operator \bowtie as a derived operator formed from \times , π , and σ . We also allow for user-defined operators to appear in expressions. The basic operators should therefore be considered as those which have “built-in” support, but they are not the only operators supported.

The basic operators differ in their behavior with respect to arity. Assume expression E_1 has arity r and expression E_2 has arity s . Then the arity of $E_1 \cup E_2$, $E_1 \cap E_2$, and $E_1 - E_2$ is $r = s$; the arity of $E_1 \times E_2$ is $r + s$; the arity of $\sigma_c(E_1)$ is r ; and the arity of $\pi_I(E_1)$ is $|I|$.

We define an additional operator which may be used in relational expressions called the *Skolem function*. A Skolem function has a name and a set of indexes. Let f be a Skolem function on indexes I . Then $f_I(E_1)$ is an expression of arity $r + 1$. Intuitively, the meaning of the operator is to add an attribute to the output, whose

values are some function f of the attribute values identified by the indexes in I . We do not provide a formal semantics here. Skolem functions are used internally as a technical device in Section 3.5.

We consider constraints of two forms. A *containment constraint* is a constraint of the form $E_1 \subseteq E_2$, where E_1 and E_2 are relational expressions. An *equality constraint* is a constraint of the form $E_1 = E_2$, where E_1 and E_2 are relational expressions. We denote sets of constraints with capital Greek letters and individual constraints with lowercase Greek letters.

A *signature* is a function from a set of relation symbols to positive integers which give their arities. In this paper, we use the terms signature and schema synonymously. We denote signatures with the letter σ . (We denote relation symbols with uppercase Roman letters R, S, T , etc.) We sometimes abuse notation and use the same symbol σ to mean simply the domain of the signature (a set of relations).

An *instance* of a database schema is a database that conforms to that schema. We use uppercase Roman letters A, B, C , etc to denote instances. If A is an instance of a database schema containing the relation symbol S , we denote by S^A the contents of the relation S in A .

Given a relational expression E and a relational symbol S , we say that E is *monotone* in S if whenever instances A and B agree on all relations except S and $S^A \subseteq S^B$, then $E(A) \subseteq E(B)$. In other words, E is monotone in S if adding more tuples to S only adds more tuples to the query result. We say that E is *anti-monotone* in S if whenever A and B agree on all relations except S and $S^A \subseteq S^B$, then $E(A) \supseteq E(B)$.

The *active domain* of an instance is the set of values that appear in the instance. We allow the use of a special relational symbol D which denotes the active domain of an instance. D can be thought of as a shorthand for the relational expression $\bigcup_{i=1}^n \bigcup_{j=1}^{a_i} \pi_j(S_i)$ where $\sigma = \{S_1, \dots, S_n\}$ is the signature of the database and a_i is the arity of relation S_i . We also allow the use of another special relation in expressions, the empty relation \emptyset .

An instance A *satisfies* a containment constraint $E_1 \subseteq E_2$ if $E_1(A) \subseteq E_2(A)$. An instance A *satisfies* an equality constraint $E_1 = E_2$ if $E_1(A) = E_2(A)$. We write $A \models \xi$ if the instance A satisfies the constraint ξ and $A \models \Sigma$ if A satisfies every constraint in Σ . Note that $A \models E_1 = E_2$ iff $A \models E_1 \subseteq E_2$ and $A \models E_2 \subseteq E_1$.

EXAMPLE 2 The constraint that the first attribute of a binary relation S is the key for the relation, which can be expressed in a logic-based setting as the equality-generating dependency $S(x, y), S(x, z) \rightarrow y = z$ may be expressed in our setting as a containment constraint by making use of the active domain relation

$$\pi_{24}(\sigma_{1=3}(S^2)) \subseteq \sigma_{1=2}(D^2)$$

where S^2 is short for $S \times S$ and D^2 is short for $D \times D$. \square

A *mapping* is a binary relation on instances of database schemas. We reserve the letter m for mappings. Given a class of constraints \mathcal{L} , we associate to every expression of the form $(\sigma_1, \sigma_2, \Sigma_{12})$ the mapping

$$\{\langle A, B \rangle : (A, B) \models \Sigma_{12}\}.$$

That is, it defines which instances of two schemas correspond to each other. Here Σ_{12} is a finite subset of \mathcal{L} over the signature $\sigma_1 \cup \sigma_2$, σ_1 is the *input (or source) signature*, σ_2 is the *output (or target) signature*, A is a database with signature σ_1 and B is a database with signature σ_2 . We assume that σ_1 and σ_2 are disjoint. (A, B) is the database with signature $\sigma_1 \cup \sigma_2$ obtained by taking all the relations in A and B together. Its active domain is the union of the

active domains of A and B . In this case, we say that m is given by $(\sigma_1, \sigma_2, \Sigma_{12})$.

Given two mappings m_{12} and m_{23} , the *composition* $m_{12} \circ m_{23}$ is the unique mapping

$$\{\langle A, C \rangle : \exists B (\langle A, B \rangle \in m_{12} \text{ and } \langle B, C \rangle \in m_{23})\}.$$

Assume two mappings m_{12} and m_{23} are given by $(\sigma_1, \sigma_2, \Sigma_{12})$ and $(\sigma_2, \sigma_3, \Sigma_{23})$. The *mapping composition problem* is to find Σ_{13} such that $m_{12} \circ m_{23}$ is given by $(\sigma_1, \sigma_3, \Sigma_{13})$.

Given a finite set of constraints Σ over some schema σ and another finite set of constraints Σ' over some subschema σ' of σ we say that Σ is *equivalent* to Σ' , denoted $\Sigma \equiv \Sigma'$, if

1. (**Soundness**) Every database A over σ satisfying Σ when restricted to only those relations in σ' yields a database A' over σ' that satisfies Σ' and
2. (**Completeness**) Every database A' over σ' satisfying the constraints Σ' can be extended to a database A over σ satisfying the constraints Σ by adding new relations in $\sigma - \sigma'$ (not limited to the domain of A').

EXAMPLE 3 The set of constraints

$$\Sigma := \{R \subseteq S, S \subseteq T\}$$

is equivalent to the set of constraints $\Sigma' := \{R \subseteq T\}$. \square

Given this definition, we can restate the composition problem as follows. Given a set of constraints Σ_{12} over $\sigma_1 \cup \sigma_2$ and a set of constraints Σ_{23} over $\sigma_2 \cup \sigma_3$, find a set of constraints Σ_{13} over $\sigma_1 \cup \sigma_3$ such that $\Sigma_{12} \cup \Sigma_{23} \equiv \Sigma_{13}$.

3. ALGORITHM

3.1 Overview

At the heart of the composition algorithm (which appears at the end of this subsection), we have the procedure ELIMINATE which takes as input a finite set of constraints Σ over some schema σ that includes the relation symbol S and which produces as output another finite set of constraints Σ' over $\sigma - \{S\}$ such that $\Sigma' \equiv \Sigma$, or reports failure to do so. On success, we say that we have *eliminated* S from Σ .

Given such a procedure ELIMINATE we have several choices on how to implement COMPOSE, which takes as input three schemas σ_1, σ_2 , and σ_3 and two sets of constraints Σ_{12} and Σ_{23} over $\sigma_1 \cup \sigma_2$ and $\sigma_2 \cup \sigma_3$ respectively. The goal of COMPOSE is to return a set of constraints Σ_{13} over $\sigma_1 \cup \sigma_3$. That is, its goal is to eliminate the relation symbols from σ_2 . Since this may not be possible, we aim at eliminating from $\Sigma := \Sigma_{12} \cup \Sigma_{23}$ a set S of relation symbols in σ_2 which is as large as possible or which is maximal under some other criterion than the number of relation symbols in it. There are many choices about how to do this, but we do not explore them in this paper. Instead we simply follow the user-specified ordering on the relation symbols in σ_2 and try to eliminate as many as possible in that order.¹

We therefore concentrate in the remainder of Section 3 on ELIMINATE. It consists of the following three steps, which we describe in more detail in following sections:

¹Note that which symbols will be eliminated will in general depend on this user-defined order. Consider, for example, the constraints in the proof of Theorem 1 in [8] duplicated for two symbols S_1 and S_2 : exactly one of them can be eliminated and this will depend on the order.

1. View unfolding
2. Left compose
3. Right compose

Each of the steps 1, 2, and 3 attempts to remove S from Σ . If any of them succeeds, ELIMINATE terminates successfully. Otherwise, ELIMINATE fails.

All three steps work in essentially the same way: given a constraint that contains S alone on one side of a constraint and an expression E on the other side, they substitute E for S in all other constraints.

EXAMPLE 4 Here are three one-line examples of how each of these three steps transforms a set of two constraints into an equivalent set with just one constraint:

1. $S = R \times T, \pi(U) - S \subseteq U \Rightarrow \pi(U) - (R \times T) \subseteq U$
2. $R \subseteq S \cap V, S \subseteq T \times U \Rightarrow R \subseteq (T \times U) \cap V$
3. $T \times U \subseteq S, S - \pi(U) \subseteq R \Rightarrow (T \times U) - \pi(U) \subseteq R \quad \square$

To perform such a substitution we need an expression that contains S alone on one side of a constraint. This holds in the example, but is typically not the case. Another key feature of our algorithm is that it performs *normalization* as necessary to put the constraints into such a form. In the case of left and right compose, we also need all other expressions that contain S to be monotone in S .

We now give a more detailed technical overview of the three steps we have just introduced. To simplify the discussion below, we take $\Sigma_0 := \Sigma$ to be the input to ELIMINATE and Σ_s to be the result after step s is complete. We use E, E_1, E_2 to stand for arbitrary relational expressions and $M(S)$ to stand for a relational expression monotonic in S .

1. **View Unfolding.** We look for a constraint ξ of the form $S = E_1$ in Σ_0 where E_1 is an arbitrary expression that does not contain S . If there is no such constraint, we set $\Sigma_1 := \Sigma_0$ and go to Step 2. Otherwise, to obtain Σ_1 we remove ξ and replace every occurrence of S in every other constraint in Σ_0 with E_1 . Then $\Sigma_1 \equiv \Sigma_0$. Soundness is obvious and to show completeness it is enough to set $S = E_1$.
2. **Left Compose.** If S appears on both sides of some constraint in Σ_1 , we exit. Otherwise, we convert every equality constraint $E_1 = E_2$ that contains S into two containment constraints $E_1 \subseteq E_2$ and $E_2 \subseteq E_1$ to obtain Σ'_1 .

Next we check Σ'_1 for *right-monotonicity* in S . That is, we check whether every expression E in which S appears to the right of a containment constraint is monotonic in S . If this check fails, we set $\Sigma_2 := \Sigma_1$ and go to Step 3.

Next we left-normalize every constraint in Σ'_1 for S to obtain Σ''_1 . That is, we replace all constraints in which S appears on the left with a single equivalent constraint ξ of the form $S \subseteq E_1$. That is, S appears alone on the left in ξ . This is not always possible; if we fail, we set $\Sigma_2 := \Sigma_1$ and go to Step 3.

If S does not appear on the left of any constraint, then we add to Σ''_1 the constraint $\xi : S \subseteq E_1$ and we set $E_1 := D^r$ where r is the arity of S . Here D is a special symbol which stands for the active domain. Clearly, any S satisfies this constraint.

Now to obtain Σ'''_1 from Σ''_1 we remove ξ and for every constraint in Σ''_1 of the form $E_2 \subseteq M(S)$ where M is monotonic in S , we put a constraint of the form $E_2 \subseteq M(E_1)$ in Σ'''_1 . We call this step *basic left-composition*. Finally, to the extent that our knowledge of the operators allows us, we attempt to

eliminate D^r (if introduced) from any constraints, to obtain Σ_2 . For example $E_1 \cap D^r$ becomes E_1 .

Then $\Sigma_2 \equiv \Sigma_1$. Soundness follows from monotonicity since $E_2 \subseteq M(S) \subseteq M(E_1)$ and to show completeness it is enough to set $S := E_1$.

3. **Right Compose.** Right compose is dual to left-compose. We check for left-monotonicity and we right-normalize as in the previous step to obtain Σ''_2 with a constraint ξ of the form $E_1 \subseteq S$. If S does not appear on the right of any constraint, then we add to Σ''_2 the constraint $\xi : E_1 \subseteq S$ and set $E_1 := \emptyset$. Clearly, any S satisfies this constraint.

In order to handle projection during the normalization step we may introduce Skolem functions. For example $R \subseteq \pi_1(S)$ where R is unary and S is binary becomes $f(R) \subseteq S$. The expression $f(R)$ is binary and denotes the result of applying some unknown Skolem function f to the expression R . The right-normalization step always succeeds for select, project, and join, but may fail for other operators. If we fail to normalize, we set $\Sigma_3 := \Sigma_2$ and we exit.

Now to obtain Σ'''_2 from Σ''_2 we remove ξ and for every constraint in Σ''_2 of the form $M(S) \subseteq E_2$ where M is monotonic in S , we put a constraint of the form $M(E_1) \subseteq E_2$ in Σ'''_2 . We call this step *basic right-composition*. Finally, to the extent that our knowledge of the operators allows us, we attempt to eliminate \emptyset (if introduced) from any constraints, to obtain Σ_3 . For example $E_1 \cup \emptyset$ becomes E_1 .

Then $\Sigma'''_2 \equiv \Sigma_2$. Soundness follows from monotonicity since $M(E_1) \subseteq M(S) \subseteq E_2$ and to show completeness it is enough to set $S := E_1$.

Since during normalization we may have introduced Skolem functions, we now need a right-denormalization step to remove such Skolem functions. Following [8], we call this part deskolemization. Deskolemization is very complex and may fail. If it does, we set $\Sigma_3 := \Sigma_2$ and we exit. Otherwise, we set Σ_3 to be the result of deskolemization.

Procedure ELIMINATE

Input: Signature σ
Constraints Σ
Relation Symbol S

Output: Constraints Σ' over σ or $\sigma - \{S\}$

1. $\Sigma' := \text{VIEWUNFOLD}(\Sigma, S)$. On success, return Σ' .
2. $\Sigma' := \text{LEFTCOMPOSE}(\Sigma, S)$. On success, return Σ' .
3. $\Sigma' := \text{RIGHTCOMPOSE}(\Sigma, S)$. On success, return Σ' .
4. Return Σ and indicate failure.

Procedure COMPOSE

Input: Signatures $\sigma_1, \sigma_2, \sigma_3$
Constraints Σ_{12}, Σ_{23}
Relation Symbol S

Output: Signature σ satisfying $\sigma_1 \cup \sigma_3 \subseteq \sigma \subseteq \sigma_1 \cup \sigma_2 \cup \sigma_3$
Constraints Σ over σ

1. Set $\sigma := \sigma_1 \cup \sigma_2 \cup \sigma_3$.
2. Set $\Sigma = \Sigma_{12} \cup \Sigma_{23}$.
3. For every relation symbol $S \in \sigma_2$ do:
 4. $\Sigma := \text{ELIMINATE}(\sigma, \Sigma, S)$
 5. On success, set $\sigma := \sigma - \{S\}$.
6. Return σ, Σ .

THEOREM 1 *Algorithm COMPOSE is correct.*

PROOF. (Sketch) Correctness follows from the soundness and completeness of the view unfolding, left-compose, and right-compose steps, and the proof of correctness of the deskolemization algorithm in [8]. \square

3.2 View Unfolding

The goal of the unfold views step is to eliminate S at an early stage by applying the technique of view unfolding. It takes as input a set of constraints Σ_0 and a symbol S to be eliminated. It produces as output an equivalent set of constraints Σ_1 with S eliminated (in the success case), or returns Σ_0 (in the failure case). The step proceeds as follows. We look for a constraint ξ of the form $S = E_1$ in Σ_0 where E_1 is an arbitrary expression that does not contain S . If there is no such constraint, we set $\Sigma_1 := \Sigma_0$ and report failure. Otherwise, to obtain Σ_1 we remove ξ and replace every occurrence of S in every other constraint in Σ_0 with E_1 . Note that S may occur in expressions that are not necessarily monotone in S , or that contain user-defined operators about which little is known. In either case, because S is defined by an equality constraint, the result is still an equivalent set of constraints. This is in contrast to left compose and right compose, which rely for correctness on the monotonicity of expressions in S when performing substitution.

EXAMPLE 5 Suppose the input constraints are given by

$$S = R_1 \times R_2, \quad \pi(R_3 - S) \subseteq T_1, \quad T_2 \subseteq T_3 - \sigma_c(S).$$

Then unfold views deletes the first constraint and substitutes $R_1 \times R_2$ for S in the second two constraints, producing

$$\pi(R_3 - (R_1 \times R_2)) \subseteq T_1, \quad T_2 \subseteq T_3 - \sigma_c(R_1 \times R_2).$$

Note that in this example, neither left compose nor right compose would succeed in eliminating S . Left compose would fail because the expression $T_3 - \sigma_c(S)$ is not monotone in S . Right compose would fail because the expression $\pi(R_3 - S)$ is not monotone in S . Therefore view unfolding does indeed give us some extra power compared to left compose and right compose alone. \square

3.3 Checking Monotonicity

The correctness of performing substitution to eliminate a symbol S in the left compose and right compose steps depends upon the left-hand side (lhs) or right-hand side (rhs) of all constraints being monotone in S . We describe here a sound but incomplete procedure MONOTONE for checking this property. MONOTONE takes as input an expression E and a symbol S . It returns ‘m’ if the expression is monotone in S , ‘a’ if the expression is anti-monotone in S , ‘i’ if the expression is independent of S (for example, because it does not contain S), and ‘u’ (unknown) if it cannot say how the expression depends on S . For example, given the expression $S \times T$ and symbol S as input, MONOTONE returns ‘m’, while given the expression $\sigma_{c_1}(S) - \sigma_{c_2}(S)$ and the symbol S , MONOTONE returns ‘u’.

The procedure is defined recursively in terms of the six basic relational operators. In the base case, the expression is a single relational symbol, in which case MONOTONE returns ‘m’ if that symbol is S , and ‘i’ otherwise. Otherwise, in the recursive case, MONOTONE first calls itself recursively on the operands of the top-level operator, then performs a simple table lookup based on the return values and the operator. For the unary expressions $\sigma(E_1)$ and $\pi(E_1)$, we have that $\text{MONOTONE}(\sigma(E_1), S) = \text{MONOTONE}(\pi(E_1), S) = \text{MONOTONE}(E_1, S)$ (in other words, σ and π do not affect the monotonicity of the expression). Otherwise, for the binary expressions $E_1 \cup E_2$, $E_1 \cap E_2$, $E_1 \times E_2$, and $E_1 - E_2$, there are sixteen cases to consider, corresponding to the possible values of $\text{MONOTONE}(E_1, S)$ and $\text{MONOTONE}(E_2, S)$.

Note that \times , \cap , and \cup all behave in the same way from the point of view of MONOTONE, that is, $\text{MONOTONE}(E_1 \cup E_2, S) = \text{MONOTONE}(E_1 \cap E_2, S) = \text{MONOTONE}(E_1 \times E_2, S)$, for all E_1, E_2 . Set difference $-$, on the other hand, behaves differently than the others.

In order to support user-defined operators in MONOTONE, we just need to know the rules regarding the monotonicity of the operator in S , given the monotonicity of its operands in S . Once these rules have been added to the appropriate tables, MONOTONE supports the user-defined operator automatically.

3.4 Left Compose

Recall from Section 3.1 that left compose consists of four main steps, once equality constraints have been converted to containment constraints. The first is to check the constraints for right-monotonicity in S , that is, to check whether every expression E in which S appears to the right of a containment constraint is monotonic in S . Section 3.3 already described the procedure for checking this. The other three steps are left normalize, basic left compose, and eliminate domain relation. In this section we describe those steps in more detail, and we give some examples to illustrate their operation.

3.4.1 Left Normalize

The goal of left normalize is to put the set of input constraints in a form such that the symbol S to be eliminated appears on the left of exactly one constraint, which is of the form $S \subseteq E_2$. We say that the constraints are then in *left normal form*. In contrast to right normalize, left normalize does not always succeed even on the basic relational operators. Nevertheless, left composition is useful because it may succeed in cases where right composition fails for other reasons. We give an example of this in Section 3.4.2.

We make use of the following identities for containment constraints in left normalize (note that here again the subscripts of the projection operator π are omitted for readability):

$$\begin{aligned} \cup : & \quad E_1 \cup E_2 \subseteq E_3 \leftrightarrow E_1 \subseteq E_3, E_2 \subseteq E_3 \\ - : & \quad E_1 - E_2 \subseteq E_3 \leftrightarrow E_1 \subseteq E_2 \cup E_3 \\ \pi : & \quad \pi(E_1) \subseteq E_2 \leftrightarrow E_1 \subseteq E_2 \times D^r \\ \sigma : & \quad \sigma_c(E_1) \subseteq E_2 \leftrightarrow E_1 \subseteq E_2 \cup (D^r - \sigma_c(D^r)) \end{aligned}$$

To each identity in the list, we associate a rewriting rule that takes a constraint of the form given by the lhs of the identity and produces an equivalent constraint or set of constraints of the form given by the rhs of the identity. For example, from the identity for σ we obtain a rule that matches a constraint of the form $\sigma_c(E_1) \subseteq E_2$ and rewrites it into equivalent constraints of the form $E_1 \subseteq E_2 \cup (D^r - \sigma_c(D^r))$. Note that there is at most one rule for each operator. So to find the rule that matches a particular expression, we need only look up the rule corresponding to the topmost operator in the expression.

We do not know of identities covering all of the basic relational operators. In particular, constraints of the following form seem to be problematic:

$$E_1 \cap E_2 \subseteq E_3, \quad E_1 \times E_2 \subseteq E_3, \quad E_1 - E_2 \subseteq E_3,$$

where in the last constraint, E_2 is the expression that contains the symbol S . For example, one might be tempted to think that the constraint $E_1 \times E_2 \subseteq E_3$ could be rewritten as $E_1 \subseteq \pi(E_3), E_2 \subseteq \pi(E_3)$. However, the following counterexample shows that this rewriting is invalid:

EXAMPLE 6 Let R, S be unary relations and let T be a binary relation. Define the instance A to be $R^A := \{1, 2\}, S^A := \{1, 2\}$,

$T^A := \{11, 22\}$. Then $A \models \{R \subseteq \pi_1(T), S \subseteq \pi_2(T)\}$, but $A \not\models \{R \times S \subseteq T\}$. \square

In addition to the basic relational operators, left normalize may be extended to handle user-defined operators by specifying a user-defined rewriting rule for each such operator.

Left normalize proceeds as follows. Let Σ_1 be the set of input constraints, and let S be the symbol to be eliminated from Σ_1 . Left normalize computes a set Σ'_1 of constraints as follows. Set $\Gamma_1 := \Sigma_1$. We loop as follows, beginning at $i = 1$. In the i th iteration, there are two cases:

1. If there is no constraint in Γ_i that contains S on the lhs in a complex expression, set Σ'_1 to be Γ_i with all the constraints containing S on the lhs collapsed into a single constraint, which has an intersection of expressions on the right. For example, $S \subseteq E_1, S \subseteq E_2$ becomes $S \subseteq E_1 \cap E_2$. If S does not appear on the lhs of any expression, we add to Σ'_1 the constraint $S \subseteq D^r$ where r is the arity of S . Finally, return success.
2. Otherwise, choose some constraint $\xi := E_1 \subseteq E_2$, where E_1 contains S . If there is no rewriting rule for the top-level operator in E_1 , set $\Sigma'_1 := \Sigma_1$ and return failure. Otherwise, set Γ_{i+1} to be the set of constraints obtained from Γ_i by replacing ξ with its rewriting, and iterate.

EXAMPLE 7 Suppose the input constraints are given by

$$R - S \subseteq T, \quad \pi(S) \subseteq U.$$

where S is the symbol to be eliminated. Then left normalization succeeds and returns the constraints

$$R \subseteq S \cup T, \quad S \subseteq U \times D^r. \quad \square$$

EXAMPLE 8 Suppose the input constraints are given by

$$R \cap S \subseteq T, \quad \pi(S) \subseteq U.$$

Then left normalization fails, because there is no rule matching the lhs of $R \cap S \subseteq T$. \square

EXAMPLE 9 Suppose the input constraints are given by

$$R \cap T \subseteq S, \quad U \subseteq \pi(S).$$

Since there is no constraint containing S on the left, left normalize adds the trivial constraint $S \subseteq D^r$, producing

$$R \cap T \subseteq S, \quad U \subseteq \pi(S), \quad S \subseteq D^r. \quad \square$$

3.4.2 Basic Left Compose

Among the constraints produced by left normalize, there is a single constraint $\xi := S \subseteq E_1$ that has S on its lhs. In basic left compose, we remove ξ from the set of constraints, and we replace every other constraint of the form $E_2 \subseteq M(S)$, where $M(S)$ is monotonic in S , with a constraint of the form $E_2 \subseteq M(E_1)$. This is easier to understand with the help of a few examples.

EXAMPLE 10 Consider the constraints from Example 7 after left normalization:

$$R \subseteq S \cup T, \quad S \subseteq U \times D^r.$$

The expression $S \cup T$ is monotone in S . Therefore, we are able to left compose to obtain

$$R \subseteq (U \times D^r) \cup T.$$

Note although the input constraints from Example 7 could just as well be put in right normal form, right compose would fail, because the expression $R - S$ is not monotone in S . Thus left compose does indeed give us some additional power. \square

EXAMPLE 11 We continue with the constraints from Example 9:

$$R \cap T \subseteq S, \quad U \subseteq \pi(S), \quad S \subseteq D^r.$$

We left compose and obtain

$$R \cap T \subseteq D^r, \quad U \subseteq \pi(D^r).$$

Note that the active domain relation D occurs in these constraints. In the next section, we explain how to eliminate it. \square

3.4.3 Eliminate Domain Relation

We have seen that left compose may produce a set of constraints containing the symbol D which represents the active domain relation. The goal of this step is to eliminate D from the constraints, to the extent that our knowledge of the operators allows, which may result in entire constraints disappearing in the process as well. We use rewriting rules derived from the following identities for the basic relational operators:

$$\begin{aligned} E_1 \cup D^r &= D^r & E_1 \cap D^r &= E_1 \\ E_1 - D^r &= \emptyset & \pi_I(D^r) &= D^{|I|} \end{aligned}$$

(We do not know of any identities applicable to cross product or selection.) In addition, the user may supply rewriting rules for user-defined operators, which we will make use of if present. The constraints are rewritten using these rules until no rule applies. At this point, D may appear alone on the rhs of some constraints. We simply delete these, since a constraint of this form is satisfied by any instance. Note that we do not always succeed in eliminating D from the constraints. However, this is acceptable, since a constraint containing D can still be checked.

EXAMPLE 12 We continue with the constraints from Example 11:

$$R \cap T \subseteq D^r, \quad U \subseteq \pi(D^r).$$

First, the domain relation rewriting rules are applied, yielding

$$R \cap T \subseteq D^r, \quad U \subseteq D^k,$$

where k is the arity of $\pi(D^r)$. Then, since both of these constraints have the domain relation alone on the rhs, we are able to simply delete them. \square

3.5 Right Compose

Recall from Section 3.1 that right compose proceeds through five main steps. The first step is to check that every expression E that appears to the left of a containment constraint is monotonic in S . The procedure for checking this was described in Section 3.3. The other four steps are right normalize, basic right compose, right-denormalize, and eliminate empty relations. In this section, we describe these steps in more detail and provide some examples.

3.5.1 Right Normalize

Right normalize is dual to left normalize. The goal of right normalize is to put the constraints in a form where S appears on the rhs of exactly one constraint, which has the form $E_1 \subseteq S$. We say that the constraints are then in *right normal form*. We make use of the following identities for containment constraints in right normalization (note that here again the subscripts of the projection operator π are omitted for readability):

$$\begin{aligned}
\cup : \quad & E_1 \subseteq E_2 \cup E_3 \leftrightarrow E_1 - E_3 \subseteq E_2 \\
& \quad \quad \quad \leftrightarrow E_1 - E_2 \subseteq E_3 \\
\cap : \quad & E_1 \subseteq E_2 \cap E_3 \leftrightarrow E_1 \subseteq E_2, E_1 \subseteq E_3 \\
\times : \quad & E_1 \subseteq E_2 \times E_3 \leftrightarrow \pi(E_1) \subseteq E_2, \pi(E_1) \subseteq E_3 \\
- : \quad & E_1 \subseteq E_2 - E_3 \leftrightarrow E_1 \subseteq E_2, E_1 \cap E_3 \subseteq \emptyset \\
\pi : \quad & E_1 \subseteq \pi(E_2) \leftrightarrow f(E_1) \subseteq E_2 \\
\sigma : \quad & E_1 \subseteq \sigma_c(E_2) \leftrightarrow E_1 \subseteq E_2, E_1 \subseteq \sigma_c(D^r)
\end{aligned}$$

As in left normalize, to each identity in the list, we associate a rewriting rule that takes a constraint of the form given by the lhs of the identity and produces an equivalent constraint or set of constraints of the form given by the rhs of the identity. For example, from the identity for σ we obtain a rule that matches constraints of the form $E_1 \subseteq \sigma_c(E_2)$ and produces the equivalent pair of constraints $E_1 \subseteq E_2$ and $E_1 \subseteq \sigma_c(D^r)$. As with left normalize, there is at most one rule for each operator. So to find the rule that matches a particular expression, we need only look up the rule corresponding to the topmost operator in the expression. In contrast to the rules used by left normalize, there is a rule in this list for each of the six basic relational operators. Therefore right normalize always succeeds when applied to constraints that use only basic relational expressions.

Just as with left normalize, user-defined operators can be supported via user-specified rewriting rules. If there is a user-defined operator that does not have a rewriting rule, then right normalize may fail in some cases.

Note that the rewriting rule for the projection operator π may introduce Skolem functions. The deskolemize step will later attempt to eliminate any Skolem functions introduced by this rule. If we have additional knowledge about key constraints for the base relations, we use this to minimize the list of attributes on which the Skolem function depends. This increases our chances of success in deskolemize.

Right normalize proceeds as follows. Let Σ_2 be the set of input constraints, and let S be the symbol to be eliminated from Σ_2 . Right normalize computes a set Σ'_2 of constraints as follows. Set $\Gamma_1 := \Sigma_2$. We loop as follows, beginning at $i = 1$. In the i th iteration, there are two cases:

1. If there is no constraint in Γ_i that contains S on the rhs in a complex expression, set Σ'_2 to be the same as Γ_i but with all the constraints containing S on the rhs collapsed into a single constraint containing a union of expressions on the left. For example, $E_1 \subseteq S, E_2 \subseteq S$ becomes $E_1 \cup E_2 \subseteq S$. If S does not appear on the rhs of any expression, we add to Σ'_2 the constraint $\emptyset \subseteq S$. Finally, return success.
2. Otherwise, choose some constraint $\xi := E_1 \subseteq E_2$, where E_2 contains S . If there is no rewriting rule corresponding to the top-level operator in E_2 , set $\Sigma'_2 := \Sigma_2$ and return failure. Otherwise, set Γ_{i+1} to be the set of constraints obtained from Γ_i by replacing ξ with its rewriting, and iterate.

EXAMPLE 13 Consider the constraints given by

$$S \times T \subseteq U, \quad T \subseteq \sigma_c(S) \times \pi(R).$$

Right normalize leaves the first constraint alone and rewrites the second constraint, producing

$$S \times T \subseteq U, \quad \pi(T) \subseteq S, \quad \pi(T) \subseteq \sigma_c(D^r), \quad \pi(T) \subseteq \pi(R).$$

Notice that rewriting stopped for the constraint $\pi(T) \subseteq \pi(R)$ immediately after it was produced, because S does not appear on its rhs. \square

EXAMPLE 14 Consider the constraints given by

$$R \subseteq \pi(S \times (T \cap U)), \quad S \subseteq \sigma_c(T).$$

Right normalize rewrites the first constraint and leaves the second constraint alone, producing

$$\pi(f(R)) \subseteq S, \quad \pi(f(R)) \subseteq T \cap U, \quad S \subseteq \sigma_c(T).$$

Note that a Skolem function f was introduced in order to handle the projection operator. After right compose, the deskolemize procedure will attempt to get rid of the Skolem function f . \square

3.5.2 Basic Right Compose

After right normalize, there is a single constraint $\xi := E_1 \subseteq S$ which has S on its rhs. In basic right compose, we remove ξ from the set of constraints, and we replace every other constraint of the form $M(S) \subseteq E_2$, where $M(S)$ is monotonic in S , with a constraint of the form $M(E_1) \subseteq E_2$. This is easier to understand with the help of a few examples.

EXAMPLE 15 Recall the constraints produced by right normalize in Example 13:

$$S \times T \subseteq U, \quad \pi(T) \subseteq S, \quad \pi(T) \subseteq \sigma_c(D^r), \quad \pi(T) \subseteq \pi(R).$$

Given those constraints as input, basic right compose produces

$$\pi(T) \times T \subseteq U, \quad \pi(T) \subseteq \sigma_c(D^r), \quad \pi(T) \subseteq \pi(R).$$

Since the constraints contain no Skolem functions, in this case we are done. \square

EXAMPLE 16 Recall the constraints produced by right normalize in Example 14:

$$\pi(f(R)) \subseteq S, \quad \pi(f(R)) \subseteq T \cap U, \quad S \subseteq \sigma_c(T).$$

Given those constraints as input, basic right compose produces

$$\pi(f(R)) \subseteq T \cap U, \quad \pi(f(R)) \subseteq \sigma_c(T).$$

Note that composition is not yet complete in this case. We will need to try to complete the process by deskolemizing the constraints to get rid of f . This process is described in the next section. \square

3.5.3 Right-Denormalize

During right-normalization, we may introduce Skolem functions in order to handle projection. For example, we transform $R \subseteq \pi_1(S)$ where R is unary and S is binary to $f_1(R) \subseteq S$. The subscript 1 indicates that f depends on position 1 of R . That is, $f_1(R)$ is a binary expression where to every value in R another value is associated by f . Thus, after basic right-composition, we may have constraints with Skolem functions in them. The semantics of such constraints is that they hold iff there exist *some* values for the Skolem functions which satisfy the constraints. The objective of the deskolemization step is to remove such Skolem functions. It is a complex 12-step procedure based on a similar procedure presented in [8].

Procedure DESKOLEMIZE(Σ)

1. Unnest
2. Check for cycles
3. Check for repeated function symbols
4. Align variables
5. Eliminate restricting atoms
6. Eliminate restricted constraints

7. Check for remaining restricted constraints
8. Check for dependencies
9. Combine dependencies
10. Remove redundant constraints
11. Replace functions with \exists -variables
12. Eliminate unnecessary \exists -variables

Here we only highlight some aspects specific to this implementation. First of all, as we already said, we use an algebra-based representation instead of a logic-based representation. A Skolem function for us is a relational operator which takes an r -ary expression and produces an expression of arity $r + 1$. Our goal at the end of step 3 is to produce expressions of the form

$$\pi \sigma f g \dots \sigma (R_1 \times R_2 \times \dots \times R_k).$$

Here

- π selects which positions will be in the final expression,
- the outer σ selects some rows based on values in the Skolem functions,
- f, g, \dots is a sequence of Skolem functions,
- the inner σ selects some rows independently of the values in the Skolem functions, and
- $(R_1 \times R_2 \times \dots \times R_k)$ is a cross product of possibly repeated base relations.

The goal of step 4 is to make sure that across all constraints, all these expressions have the same arity for the part after the Skolem functions. This is achieved by possibly padding with the D symbol. Furthermore, step 4 aligns the Skolem functions in such a way that across all constraints the same Skolem functions appear, in the same sequence. For example, if we have the two expressions $f_1(R)$ and $g_1(S)$ with R, S unary, step 4 rewrites them as

$$\pi_{13} g_2 f_1 (R \times S) \text{ and } \pi_{24} g_2 f_1 (R \times S).$$

Here $R \times S$ is a binary expression with R in position 1 and S in position 2 and $g_2 f_1 (R \times S)$ is an expression of arity 4 with R in position 1, S in position 2, f in position 3 depending only on position 1, and g in position 4 depending only on position 2.

The goal of step 5 is to eliminate the outer selection σ and of step 6 to eliminate constraints having such an outer selection. The remaining steps correspond closely to the logic-based approach.

Deskolemization is complex and may fail at several of the steps above. The following two examples illustrate some cases where deskolemization fails.

EXAMPLE 17 Consider the following constraints from [5] where E, F, C, D are binary and $\sigma_2 = \{F, C\}$:

$$E \subseteq F, \quad \pi_1(E) \subseteq \pi_1(C), \quad \pi_2(E) \subseteq \pi_1(C) \\ \pi_{46} \sigma_{1=3,2=5} (F \times C \times C) \subseteq D$$

Right-composition succeeds at eliminating F to get

$$\pi_1(E) \subseteq \pi_1(C), \quad \pi_2(E) \subseteq \pi_1(C) \\ \pi_{46} \sigma_{1=3,2=5} (E \times C \times C) \subseteq D$$

Right-normalization for C yields

$$\pi_{13} f_{12}(E) \subseteq C, \quad \pi_{23} g_{12}(E) \subseteq C \\ \pi_{46} \sigma_{1=3,2=5} (E \times C \times C) \subseteq D$$

and basic right-composition yields 4 constraints including

$$\pi_{46} \sigma_{1=3,2=5} (E \times (\pi_{13} f_{12}(E)) \times (\pi_{13} f_{12}(E))) \subseteq D$$

which causes deskolemize to fail at step 3. Therefore, right-compose fails to eliminate C . As shown in [5] eliminating C is impossible by any means. \square

3.5.4 Eliminate Empty Relation

Right compose may introduce the empty relation symbol \emptyset into the constraints during composition in the case where S does not appear on the rhs of any constraint. In this step, we attempt to eliminate the symbol, to the extent that our knowledge of the operators allows. This is usually possible and often results in constraints disappearing entirely, as we shall see. For the basic relational operators, we make use of rewriting rules derived from the following identities for the basic relational operators:

$$E_1 \cup \emptyset = E_1 \quad E_1 \cap \emptyset = \emptyset \quad E_1 - \emptyset = E_1 \\ \emptyset - E_1 = \emptyset \quad \sigma_c(\emptyset) = \emptyset \quad \pi_I(\emptyset) = \emptyset$$

In addition we allow the user to supply rewriting rules for user-defined operators. The constraints are rewritten using these rules until no rule applies. At this point, some constraints may have the form $\emptyset \subseteq E_2$. These constraints are then simply deleted, since they are satisfied by any instance. We do not always succeed in eliminating the empty relation symbol from the constraints. However, this is acceptable, since a constraint containing \emptyset can still be checked.

4. EXPERIMENTS

Prior work [5, 8] showed that characterizing the class of mappings that can be composed is very difficult, even when no outer joins, unions, or difference operators are present in the mapping constraints. In this section we make a first attempt to explore experimentally the boundaries of mappings that can be composed. Ultimately, our goal is to develop a mapping composition benchmark that can be used to compare competing implementations of mapping composition. The experiments that we report on could form a part of such a benchmark. Due to space limitations we focus on the key results; a more detailed discussion is in [2].

All composition problems used in our experiments are available for online download in a machine-readable format.² We designed a plain-text syntax for specifying mapping composition tasks. Mapping constraints are encoded according to the index-based algebraic notation introduced in Section 2. We built a parser that takes as input a textual specification of a composition problem and converts it into an internal algebraic representation, which is fed to our algorithm. A sample run of the algorithm on mappings that contain relational operators union, difference, and left outerjoin is shown in [2].

We used two data sets in our experiments. The first one contains 22 composition problems drawn from the recent literature [5, 7, 8], which illustrate subtle composition issues. The attractiveness of this data set is that the expected output mappings were verified manually and documented in the literature, sometimes using formal proofs. So, this data set serves as a test suite that can be used for verifying implementations of composition.

The second data set used in the experiments is synthetic. Using synthetic mappings appears to be the primary way to study the scalability and effectiveness of mapping composition; this approach was also followed in [14]. Our study is based on the schema evolution scenarios outlined in the introduction. Specifically, we focus on schema editing and schema reconciliation tasks, since mapping adaptation can be viewed as a special case of schema reconciliation where one of the input mappings is fixed.

In the study based on synthetic mappings, we determine the success rate of our algorithm in eliminating symbols for various composition tasks, measure its execution time, and investigate the contributions of the main steps of the algorithm on its overall performance. To generate synthetic data for our experiments,

²<http://research.microsoft.com/db/ModelMgt/composition/>

Primitive	Description	Input relation	Output relation(s)	Mapping constraint(s)
AR	Add relation	\emptyset	$R(\underline{\mathbf{A}}, \mathbf{B})$	\emptyset
DR	Drop relation	$R(\mathbf{A})$	\emptyset	\emptyset
AA	Add attribute	$R(\mathbf{A})$	$S(\mathbf{A}, C)$	$R = \pi_{\mathbf{A}}(S)$
DA	Drop attribute	$R(\mathbf{A}), C \in \mathbf{A}$	$S(\mathbf{A} - \{C\})$	$\pi_{\mathbf{A} - \{C\}}(R) = S$
D	Add default	$R(\mathbf{A})$	$S(\mathbf{A}, C)$	$R \times \{c\} = S; R = \pi_{\mathbf{A}}(\sigma_{C=c}(S))$
H	Horizontal partitioning	$R(\mathbf{A}), C \in \mathbf{A}$	$S(\mathbf{A}), T(\mathbf{A})$	$\sigma_{C=c_S}(R) = S; \sigma_{C=c_T}(R) = T; R = S \cup T$
V	Vertical partitioning	$R(\underline{\mathbf{A}}, \mathbf{B}, \mathbf{C})$	$S(\underline{\mathbf{A}}, \mathbf{B}), T(\underline{\mathbf{A}}, \mathbf{C})$	$\pi_{\mathbf{A}, \mathbf{B}}(R) = S; \pi_{\mathbf{A}, \mathbf{C}}(R) = T; R = S \bowtie_{\mathbf{A}} T$
N	Normalization	$R(\mathbf{A}, \mathbf{B}, \mathbf{C})$	$S(\underline{\mathbf{A}}, \mathbf{B}), T(\mathbf{A}, \mathbf{C})$	same as vertical; $\pi_{\mathbf{A}}(T) \subseteq \pi_{\mathbf{A}}(S)$
Sub	Subset	$R(\mathbf{A})$	$S(\mathbf{A})$	$R \subseteq S$
Sup	Superset	$R(\mathbf{A})$	$S(\mathbf{A})$	$R \supseteq S$

Figure 1: Schema evolution primitives

we developed a tool which we call a *schema evolution simulator*. The simulator is driven by a weighted set of schema evolution primitives, such as adding or dropping attributes and relations, schema normalization, and vertical partitioning. It produces an evolved schema and a mapping between the original schema and the evolved schema.

In the schema editing scenario, we run the simulator to mimic the schema transformation operations performed by a database designer. The mapping between the original schema and the current state of the schema is composed with the mapping produced by each subsequent schema evolution primitive. We record the success or failure of each composition operation for the applied primitives.

To study schema reconciliation, we use the simulator to produce a large number of evolved schemas and mappings for a given original schema. We then compose the generated mappings pairwise using our composition tool.

Our key observations from this study are summarized below:

- Our algorithm is very effective in performing composition. In most settings, it is able to eliminate as much as a half of the symbols from the intermediate signature, and often all of them (Figures 2,5–7).
- The algorithm’s median running time is in the subsecond range for mappings containing hundreds of constraints (Figures 3–5,7).
- Composition becomes increasingly harder as more schema evolution primitives are applied to schemas (Figure 7).
- Certain kinds of schema evolution primitives are more likely to produce complications for composition than others (Figures 2,3,5).
- Key constraints do not substantially affect the symbol-eliminating power of the algorithm, yet significantly increase the running time (Figures 2,3).
- It is beneficial to keep the symbols that could not be eliminated in the mappings as second-order constraints as long as possible. Subsequent composition operations may eliminate up to a third of those symbols [2].
- Our algorithm appears to be order-invariant on the studied data sets, i.e., it eliminates the same fraction of symbols no matter in what order the symbols are tried (in the loop in Line 3 of procedure COMPOSE in Section 3.1).

We found that the output constraints produced by our algorithm are often more verbose than the ones derived manually, so simplification of output mappings is essential. An example of such simplification is detecting and removing implied constraints. Mapping simplification appears to be a problem of independent interest and is out of scope of this paper. Furthermore, we found that our technique of representing key constraints using the active domain

symbol works well in many cases, but fails in others due to de-Skolemization. Leveraging key constraints in a more direct way, e.g., using specialized rules in the composition algorithm, may increase its coverage.

In the remaining space we present some details of our study. The experiments were conducted on a 1.5 GHz Pentium M machine.

4.1 Schema Evolution Simulator

Figure 1 lists the schema evolution primitives implemented in our simulator. Each primitive takes zero or one relation as input and produces zero or more new relations and constraints. The produced constraints link the output relations to the input relations or represent key or inclusion constraints on the output relations. In the figure, we use the named perspective on the relational algebra to simplify the exposition. Attribute lists are shown in bold (e.g., \mathbf{A}), keys are underlined, and lower case symbols denote constants. The shown list of primitives covers a large fraction of those used in the schema evolution and data integration literature.

The first four primitives AR, DR, AA, and DA add or delete relations and attributes. Primitive AR creates a new relation. The arity of the new relation is chosen at random between some minimal and maximal relation arity (2 and 10, in our study). If keys are enabled, the created relation may have a key whose size is chosen between some minimal and maximal value (1 and 3, in our study). Primitive AA adds a new attribute C to the relation R , i.e., produces a new relation S that contains C and all existing attributes \mathbf{A} of R . The mapping constraint $R = \pi_{\mathbf{A}}(S)$ states that R can be reconstructed as a projection on S . Primitive DA is complementary to AA.

Primitives D, H, V, and N have forward and backward variants. The forward variant, labeled with subscript ‘f’, contains only the constraints that define the output relations in terms of the input relations. The backward variant (‘b’) contains only the constraints that define the inputs in terms of the outputs (all forward and backward variants are listed on the x-axis in Figure 2).

The forward and backward variants reflect distinct evolution scenarios, as we explain in more detail in [2]. For example, primitive D adds an attribute C with a default value c . The forward variant D_f outputs the mapping constraint $R \times \{c\} = S$, while D_b outputs $R = \pi_{\mathbf{A}}(\sigma_{C=c}(S))$. Thus, D_f states that S is determined by R and that the newly added attribute C contains c -values only. D_b produces a weaker constraint that allows attribute $S.C$ to contain other values beyond c .

Primitive H performs a lossless horizontal partitioning of the input relation R . The vertical partitioning primitives V, V_f, V_b are the only ones that require the input relation R to have a key. The attribute set of R is partitioned across output relations S and T . Primitive N captures the schema normalization rule from database textbooks. The constants c, c_T , and c_S used in the primitives D and

