WPI-CS-TR-00-25 December 2000

Incremental Maintenance of Schema-Restructuring Views in *SchemaSQL*

by

Andreas Koeller
Elke A. Rundensteiner

Computer Science
Technical Report
Series

WORCESTER POLYTECHNIC INSTITUTE

Computer Science Department
100 Institute Road, Worcester, Massachusetts 01609-2280

Incremental Maintenance of Schema-Restructuring Views in *SchemaSQL*

Andreas Koeller and Elke A. Rundensteiner
Computer Science Department, Worcester Polytechnic Institute
{koeller, rundenst}@cs.wpi.edu

January 15, 2001

Abstract

The integration of data, especially from heterogeneous data sources (ISs), is a hard and widely studied problem. One particularly challenging issue is the integration of sources that are semantically equivalent (i.e., whose states can be mapped onto each other by an isomorphism) but schematically heterogeneous. While two such data sources may represent the same information, one may store the information inside tuples (*data*) while the other may store it in attribute or relation names (*schema*). The *SchemaSQL* query language is a recent solution to this problem powerful enough to restructure such sources into each other without the loss of information. However, the issue of maintenance of *SchemaSQL* views, once materialized over semantically heterogeneous sources, has not yet been addressed. In this paper, we propose an incremental view maintenance strategy for such schema-restructuring views. Our strategy based on an algebraic representation of the view query correctly transforms a data update or a schema change to a source into data updates, schema changes, and even mixed sequences of schema and data update requests to be applied to the view. We give a proof of the correctness of the strategy. We have also developed a prototype implementation of both a *SchemaSQL* query processor and a *SchemaSQL* View Maintainer, and then used it to compare the performance of incremental view maintenance versus complete view recomputation for this new class of schema-restructuring views. We describe both the implementation and the experiments in the paper and conclude that in many cases incremental view maintenance in *SchemaSQL* is significantly faster than recomputation.

Keywords: Heterogeneous Databases, Materialized Views, SchemaSQL, Incremental View Maintenance, Schema Restructuring.

1 Introduction

Information sources, especially on the Web, are increasingly independent from each other, being designed, administered and maintained by a multitude of autonomous data providers. Issues in data integration include the heterogeneity of data and query models across different sources, called model heterogeneity [FRV95, GRVB98, HGMI+95] and incompatibilities in schematic representations of different sources even when using the same data model, called schema heterogeneity [MIR93, LSS96]. Overcoming these problems is critical in achieving integration of a wide variety of information sources. Much work on these problems has dealt with the integration of schematically different sources under the assumption that all "data" is stored in tuples and all "schema" is stored in attributes. We will focus on another aspect of this issue in our paper, namely on the integration of heterogeneous sources under the assumption that schema elements may express data and vice versa.

One recent promising approach at overcoming such schematic heterogeneity is the language *SchemaSQL*, an SQL-extension devised by Lakshmanan et al. [LSS96, LSS99]. *SchemaSQL* allows to soften the distinction between schema and data in the relational data model by allowing to query schema (such as lists of attribute or relation names) in SQL-queries and also to use sets of values obtained from *data tuples* as *schema* in the output relation. This concept leads to a versatile query language which among other features allows to transform semantically equivalent but syntactically different schemas [LSS96, Mil98] into each other. Similar to SQL-views, *SchemaSQL*-views can be used to transform relational databases into whatever format is required by a (relational) data integration system. Therefore, *SchemaSQL* makes it possible to include a larger class of information sources into an information system.

However, the issue of view maintenance in such a system is still open. View maintenance in a *SchemaSQL* view is non-trivial, especially since such views can transform data into schema and vice versa as illustrated below. In this paper, we present the first incremental maintenance strategy for *SchemaSQL*. The strategy works correctly not only under data updates, but also under schema changes.

1.1 Motivating Example

Fig. 1 gives an example of a *SchemaSQL* query to demonstrate the capabilities of this language. Note that the two relational schemas in Fig. 1 are able to hold the same information and can be mapped into each other using *SchemaSQL* queries. The view query restructures the input relations on the left side representing

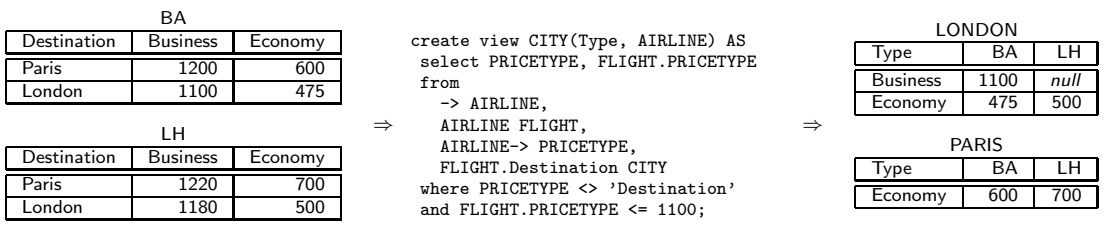

Figure 1: A *SchemaSQL* Query and its Effects.

airlines into attributes of the output relations on the right side representing destinations. The *arrow*-operator (`->`) attached to an element in the FROM-clause of a *SchemaSQL*-query allows to query schema elements, giving *SchemaSQL* its meta-data restructuring power. Standing by itself, it refers to "all relation names in that database", attached to a relation name it means "all attribute names in that relation".

SchemaSQL is also able to transform data into schema. For example, *data* from the attribute Destination in the input schema is transformed into *relation names* in the output schema, and vice versa *attribute names* in the input (Business and Economy) are restructured into *data*.

Now consider an update to one of the base relations in our example. Let a tuple $t(\mathsf{Destination} \Rightarrow \mathsf{Berlin}, \mathsf{Business} \Rightarrow 1400, \mathsf{Economy} \Rightarrow 610)$ be added to the base table LH (a **data update**). The change to the output would be the addition of a new relation Berlin (a **schema change**) with the same schema as the other two relations. This new relation would contain one tuple $t(\mathsf{Type} \Rightarrow \mathsf{Economy}, \mathsf{BA} \Rightarrow \mathit{null}, \mathsf{LH} \Rightarrow 610)$. In this example, a data update is transformed into a schema change, but all other combinations are also possible. The effect of the propagation of an update in such a query depends on numerous factors, such as the input schema, the view definition, the set of unique values in the attribute Destination across *all* input relations (city names), and the set of input relations (airline codes). For example, if a value Berlin already existed in one of the input tables, the propagation would also depend on whether other airlines offer a flight to Berlin in the Economy-class.

In summary, a schema-restructuring view must be able to propagate arbitrary data updates and schema changes that occur in its input. In the process of propagation, any such type of change could be transformed into any other in the view. This situation is significantly more complicated than update propagation in the common SQL views [BLT86, QW91, GMS93, GL95, MKK97, et al.] which deals only with propagation of data updates that are never transformed into schema changes. This problem, to our knowledge, has not been studied before.

1.2 Contributions

In the context of schema-restructuring views, there are several new issues that we must address. First, it is not sufficient to consider data updates (DUs) for *SchemaSQL*, but also schema changes (SCs). Also, *SchemaSQL* views can transform schema into data and vice-versa, thus requiring a framework that can propagate $\{DU|SC\}^* \rightarrow \{DU|SC\}^*$. As shown in this paper, using the standard approach of generating query expressions that compute some kind of "delta" relation Δ between the old and the new view after an update is not possible for *SchemaSQL*, since the schema of Δ would not be defined.

The contributions of this paper are as follows: (1) we give an algebra-based solution to the problem of incremental view maintenance of schema-restructuring views defined in *SchemaSQL*, (2) we prove this approach correct by a method similar to the equational reasoning given in [GL95], (3) we present a prototype implementation of a query processor for a subset of *SchemaSQL* and an incremental view maintenance system in Java over JDBC-capable databases, and (4) we describe experiments we have conducted on our implementation to gain insights into the performance of our algorithm.

1.3 Outline of Paper

Section 2 reviews some background on *SchemaSQL*, in particular the additional algebra operators used in *SchemaSQL* evaluation, Section 3 explains our view maintenance strategy and Section 4 proves correctness of our approach. Section 5 gives a brief overview over our implementation. Section 6 shows the results of our performance experiments. Finally, Sections 7 and 8 give related work and conclusions, respectively.

2 Background

2.1 Notation

A *value* is an element of data that is stored in a relation. Examples include strings, numbers, and dates. A *domain* D is a set of *values*.[1] D_N is the special domain of "attribute- and relation names". We implicitly assume that there is a bijective mapping from some domains to D_N. This means that the values of some, but not necessarily all, domains can be converted to names and vice versa.

A *relation* is a 3-tuple $R = (n, S, E)$ with $n \in D_N$ (the relation name), $S = (a_1, a_2, \ldots, a_n) \in (D_N)^n$ (the schema—a tuple of n attribute names) and $E \subseteq \{D_1 \times D_2 \times \ldots \times D_n\}$ (the relation extent, which is a subset of the cross-product of the domains $D_1 \times D_2 \times \ldots \times D_n$). Note that this definition associates exactly one value in S with each domain from which E is constructed (the name of an "attribute").

A *relational tuple* $t \in E$ is an n-tuple and is an element of a relation's extent. An operator $t[a_{l_1}, a_{l_2}, \ldots, a_{l_k}]$ returns the projection of t on the attributes named $a_{l_1}, a_{l_2}, \ldots, a_{l_k}$. We also define $t[*\backslash\{a_1, \ldots, a_n\}]$ to be the projection of t onto all its attributes except the ones named $a_1, \ldots, a_n$.

An *attribute* $A_i \subseteq D_i$ is a multiset that is constructed as follows: $A_i = \{t[a_i] \mid t \in E\}$, or short $A_i = E[a_i]$. Then *attribute* A_i has *attribute name* a_i. Note that we denote *attributes* by capital letters (as they are sets) and *attribute names* by small letters. We extend this notation for $E[a_1, \ldots, a_k]$ to mean the *projection* of extent E on the attributes $A_1, \ldots, A_k$. For readability, if we refer to attribute A of $R = (n, S, E)$, we actually mean the pair $A = (a, E[a])$, with $a \in S$. The term *prime attribute* refers to an attribute that is a member of any key of R and the term *non-prime attribute* refers to an attribute that is not a member of any key of R (cf. [Ull89]). The *distinct*-operator $\langle a_i \rangle$ on an attribute A_i in extent E returns the set of distinct values in A_i by removing all duplicates from the multiset $E[a_i]$.

Functional dependencies in R are defined as usual (cf. [Ull89, Chapter 7]), with $X \rightarrow A$ defining the attribute A to be functionally dependent on the set of attributes X (i.e., for any $t \in E$, the value of $t[a]$ depends only on $t[x_1, \ldots, x_k]$) . Likewise, we assume the usual definitions of *natural join* $\bowtie$ and *cross product* $\times$.

2.2 *SchemaSQL*

In relational databases it is possible to store equivalent data in different schemas, as Miller et al. [MIR93] have shown. It is also possible under certain conditions to transform data in such schemas into each other without the loss of information [Mil98]. *SchemaSQL* is an SQL derivative designed by Lakshmanan et al. [LSS96] which can be used to achieve such schema transformations.

In [LSS99], Lakshmanan et al. describe an extended algebra and algebra execution strategies to implement a *SchemaSQL* query evaluation system. It extends the standard SQL algebra which uses operators such as $\sigma(R)$, $\pi(R)$, and $R \bowtie S$ by adding four operators named UNITE, FOLD, UNFOLD, and SPLIT originally introduced by Gyssens et al. [GLS96]. Lakshmanan et al. show that any *SchemaSQL* query can be translated into this extended algebra.

We now define the four operators used in SchemaSQL in a concise manner. Examples for the four operators defined in this section can be found in Fig. 2. We will refer to the input relation of each operator as R and to the output relation as Q.

[1] Throughout this paper, we will use capital letters R to denote *(multi)sets* and small letters a to denote elements of sets.

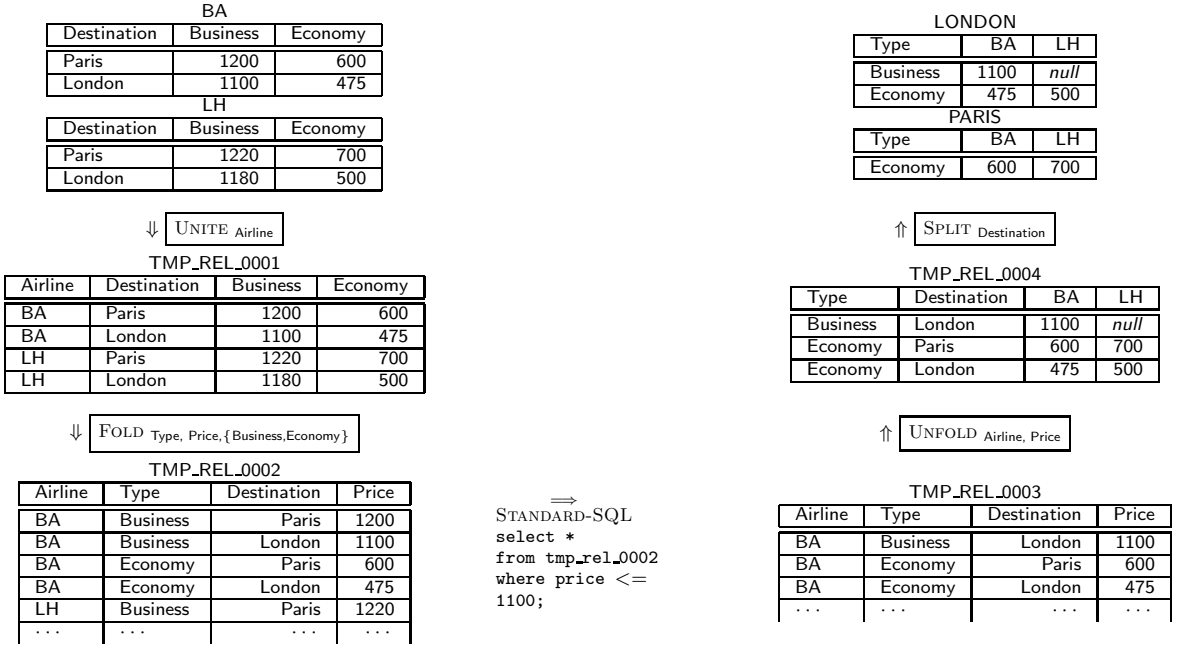

Figure 2: The Four *SchemaSQL* Operators UNITE, FOLD, UNFOLD, SPLIT.

2.2.1 UNITE

This operator is defined on a set of k relations $R^* = \{R_1, \ldots, R_k\}$ with a_p as an argument. We define a set $N^* = \{n_{R_1}, \ldots, n_{R_k}\}$, which is the set of all relation names in R^*, in the new domain D_p. We further denote by N_k the relation $R(n, E, S)$ with $N = n_k$. Then, for each R_i, we assume $S_{R_i} = (a_1, \ldots, a_n)$ and $E_{R_i} \subseteq \{D_1 \times \ldots \times D_n\}$. Note that this implies that all R_i have the same schema. The output of the UNITE operator is then one relation $Q = \text{UNITE}_{a_p}(R^*)$ with:

$E_Q \subseteq \{D_1 \times \ldots \times D_n \times D_p\}$ and $S_Q = (a_1, \ldots, a_n, a_p)$ with $E_Q = \bigcup_{n_k \in N^*} (N_k \times \{n_k\})$. Note that N_k appears both as a *relation* (N_k) and a *relation name* (n_k). In words, a new relation is constructed by taking the union of all input relations and adding a new attribute A_p whose values are the *relation names* of the input relations.

In Fig. 2, the UNITE-operator is defined over the set of relations BA, LH and has the attribute name Airline as its argument.

2.2.2 FOLD

The FOLD-operator works on a relation $R = (n_R, E_R, S_R)$ with

$$E_R \subseteq \{D_1 \times \ldots \times D_n \times \underbrace{D_d \times \ldots \times D_d}_{k \text{ times}}\}$$

and $S_R = (a_1, \ldots, a_n, a_{n+1}, \ldots a_{n+k})$ and takes as arguments the *names* of the pivot and data attributes a_p and a_d in its output relation. Note that this definition requires that k attributes of R have to be of the same domain. With R having $n + k$ attributes, we then define:

$Q = \text{FOLD}_{a_p, a_d}(R) = (n_Q, E_Q, S_Q)$ with $n_Q = n_R$, $E_Q \subseteq \{D_1 \times \ldots \times D_n \times D_d \times D_p\}$ and $S_Q =$

$(a_1, \ldots, a_n, a_d, a_p)$. We define $A^* = \{a_{n+1}, \ldots a_{n+k}\}$ as a set of values in a new domain D_p where the values are obtained by the above-mentioned conversion of attribute names into data values. Finally, $E = \bigcup_{a_k \in A^*} (R[a_1, \ldots, a_n, a_k] \times \{a_k\})$. In words, the operator takes all data values from the set of related attributes, and sorts them into *one* new attribute a_d, introducing another new attribute a_p that holds the former attribute names. Note that, since a_p becomes part of a key for a_d, it has to be included in the set X for any functional dependency $X \to A_d$.

In Fig. 2, the FOLD-operator is defined on relation TMP_REL_0001 and has the arguments $a_p =$ Type, $a_d =$ Price, $A^* =$ {Business,Economy}.

2.2.3 UNFOLD

The UNFOLD-operator is the inverse of FOLD.

The UNFOLD-operator on a relation $R = (n_R, E_R, S_R)$ with $E_R \subseteq \{D_1 \times \ldots \times D_n \times D_p \times D_d\}$ and $S_R = (a_1, \ldots, a_n, a_p, a_d)$ takes two arguments a_p, a_d which are attribute names from S_R. To simplify the notation and without loss of generality, we reorder the attributes in R (by exchanging the indices on both S_R and E_R accordingly), such that A_p and A_d become the last two attributes in R. We call A_p the *pivot attribute* and A_d the *data attribute*. Let R have $n + 2$ attributes. Let us further impose two conditions on functional dependencies in R: $(X \to Y) \Rightarrow A_d \notin X$ and $\exists (X \to A_d)$ with $A_p \in X$. That is, A_d must be non-prime, A_p must be prime and A_d must depend on A_p. We also set $A^* = R\langle A_p \rangle$ (the set of distinct values in A_p), $k = |A^*|$ and impose a total order on A^* to assign an index $1 \leq i \leq k$ to each of its elements $(A^* = \{a_1^*, \ldots, a_k^*\})$.

Then, $Q = \text{UNFOLD}_{a_p, a_d}(R) = (n_Q, S_Q, E_Q)$ with $N_Q = n_R$,

$$E_Q \subseteq \{D_1 \times \ldots \times D_n \times \underbrace{D_d \times \ldots \times D_d}_{k \text{ times}}\}$$

and $S_Q = (a_1, \ldots, a_n, a_1^*, \ldots, a_k^*)$. The extent is constructed by $E_Q = E_R[a_1, \ldots, a_n] \bowtie E_1 \bowtie \ldots \bowtie E_k$ with $E_i = \{t[a_1, \ldots, a_n, a_d] \mid t \in E_R \ \wedge \ t[a_p] = a_i^*\}$.

In words, the schema of Q consists of all attributes in R except the data and pivot attribute, plus one attribute for each distinct data value in the pivot attribute. Then, each tuple t' in Q is constructed by taking a tuple t in R and filling each new attribute A_i with the value from attribute A_d in a tuple from R that has the *name* a_i as *value* in A_p (assuming an implicit conversion between names and values as required above). Note that this mapping produces 1NF-output relations only if A_p is a prime attribute. The new attributes all have the domain D_d of the old attribute A_d.

In Fig. 2, the UNFOLD-operator is defined over relation TMP_REL_0003 and takes as its arguments $a_p =$ Airline and $a_d =$ Price. The operator then produces output by taking tuples from TMP_REL_0003, and filling the attributes representing airlines with values from the data attribute Price in TMP_REL_0003, matching attribute names in the output relation with the values of the pivot attribute Airline in the input relation. Note that this mapping produces 1NF-output relations only if A_p is a prime attribute. The new attributes all have the domain D_d of the old attribute A_d.

2.2.4 SPLIT

The SPLIT-operator is the inverse of the UNITE-operator. It transforms a single relation $R = (n_R, E_R, S_R)$ with $E_R \subseteq \{D_1 \times \ldots \times D_n \times D_p\}$ and $S_R = (a_1, \ldots, a_n, a_p)$ into a *set* of k relations with the same schema.

It takes as argument the name of the pivot attribute a_p which we assume to be the last in R. We require that A_p does not have NULL-values, i.e., $\forall x \in A_p : x \neq \perp$. The output of SPLIT is a set of relations $Q^* = \text{SPLIT}_{a_p}(R) = \{Q_1, \ldots, Q_k\}$ with $A^* = R\langle A_p \rangle$ and $k = |A^*|$. We will refer to the ordered elements of A^* as in the UNFOLD-case, i.e., $A^* = \{a_1^*, \ldots, a_k^*\}$. For each output relation Q_i, we have:

$n_{Q_i} = a_i^*$, $E_{Q_i} \subseteq \{D_1 \times \ldots \times D_n\}$, $S_{Q_i} = (a_1, \ldots, a_n)$, and $E_{Q_i} = \{t[a_1, \ldots, a_n] \mid t \in R \wedge t[a_p] = a_i^*\}$. In words, we break down R into k relations of the same schema, with the new relation names the k distinct values from R's attribute A_p.

In Fig. 2, the SPLIT-operator is defined over relation TMP_REL_0004, takes as its only argument $a_p =$ Destination, and produces 2 tables names LONDON and PARIS.

2.2.5 Flat Schemas

As suggested by the example in Fig. 1, existing relational schemas are often built with data-carrying attribute- and relation-labels, which is an important reason for schematic heterogeneity [LSS96]. In the relational model, relation and attribute labels are assumed to not contribute to the semantic content of the relation—a principle that is often misunderstood in database design. Lakshmanan et al. show that, while there are many representations of a real-world concept in the relational model, only some such schemas will not carry semantic information in their attribute and relation names. They give a definition for the notion of such *flat schemas* which we will only use in an informal way in this paper.

Definition 1 (Flat Schema) *Assume infinite pairwise disjoint sets of names $\mathcal{N}$ and values $\mathcal{V}$. Let* $\text{dom} : \mathcal{N} \rightarrow 2^{\mathcal{V}}$ *be a partial function such that for each $n \in \mathcal{N}$, whenever* $\text{dom}(n)$ *is defined, it associates name n with a non-empty set of values* $\text{dom}(n) \subset \mathcal{V}$. *Then, a relation schema $R(A_1, \ldots, A_n)$ is said to be* flat *iff all the entries $R, A_1, \ldots, A_n$ are names. A database schema is flat if all relation schemas in it have this property.*

With the notion of schema equivalence presented in [MIR93], "flattening" a relation refers to the process of transforming a relation into its flat schema-equivalent relation, preserving the information capacity of the relation.

1. If the input schema consists of $n > 1$ tables of the same schema, apply a UNITE-operator with the name of the "pivot"-attribute as a parameter to obtain relation R' (i.e., $S \stackrel{\text{UNITE}}{\longrightarrow} R'$).

2. If R or R' contains a set of attributes $\{u_1, \ldots, u_n\}$ whose labels are not from $\mathcal{N}$, apply a fold operator on $\{u_1, \ldots, u_n\}$, with the names of the two resulting attributes B, C as parameters (i.e., $R' \stackrel{\text{FOLD}}{\longrightarrow} R''$).

2.2.6 *SchemaSQL* Query Evaluation

Similar to traditional SQL evaluation, [LSS99] proposes a strategy for *SchemaSQL* query evaluation that first constructs and then processes an algebra query tree. In that way, *SchemaSQL* can be efficiently implemented over an SQL database system, which Lakshmanan et al. have shown in [LSS99].

For simplicity of notation, we will treat a database consisting of n relations $R_1, \ldots, R_n$ of the same schema as a relation and denote it by R.

In order to evaluate a *SchemaSQL* query, an algebra expression using standard relational algebra plus the four operators introduced above is constructed. This expression is of the following form [LSS99]:

$$V = \text{SPLIT}_a(\text{UNFOLD}_{b,c}(\pi_{\bar{d}}(\sigma_{cond}(\text{FOLD}_{e_1,f_1,\bar{g_1}}(\text{UNITE}_h(R_1)) \times \ldots \times \text{FOLD}_{e_m,f_m,\bar{g_m}}(\text{UNITE}_h(R_m)))))) \quad (1)$$

with attribute names a, b, c, e_i, f_i, h_i, the sets of attribute names $\bar{d}$ and $\bar{g_i}$, and selection predicates *cond* determined by the query. Any of the four *SchemaSQL* operators may be missing from the expression (i.e., may not be needed for a particular query). $R_1 \ldots R_m$ are base relations, or, in the case that the expression contains a UNITE-operator, sets of relations with equal schema.

The algebraic expression for our running example (Fig. 1) is:

$$\begin{aligned} V \quad = \quad & \text{SPLIT}_{\text{Destination}}(\text{UNFOLD}_{\text{Airline,Price}}(\sigma_{\text{Price}<1100}(\\ & \text{FOLD}_{\text{Type, Price, \{Business,Economy\}}}(\text{UNITE}_{\text{Airline}}(\text{BA,LH}))))) \end{aligned} \quad (2)$$

This algebraic expression is then used to construct an algebra tree (Fig. 3) whose nodes are any of the four *SchemaSQL* operators or a "Standard-SQL"-operator (including the π, σ, and $\times$-operators of the algebra expression) with standard relations "traveling" along its edges. The query is then evaluated by traversing the algebra tree and executing a query processing strategy for each operator, analogous to traditional SQL query evaluation.

Note that the query tree could include $\times$-operators (which do not exist in our example), but that the order of UNITE, FOLD, UNFOLD, SPLIT (if they exist) is fixed by the template in Equation 1. The UNITE operator takes a number of relations of the same schema as an input, while the SPLIT-operator produces as output a set of relations of the same schema. Note that the algebra tree in Fig. 3 is very simple, in more complex queries, the tree could "fork" at the *Standard-SQL-node*, and several smaller "flattening" trees using UNITE- and FOLD-operators could occur. In that case, and also in the case of standard relational joins, the *Standard-SQL-node* would itself contain a more complex algebra tree containing simple SQL algebra nodes.

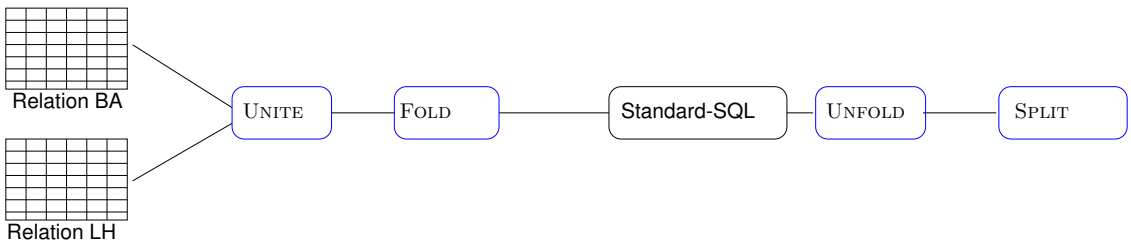

Figure 3: The Algebra Tree for the Example in Fig.1

3 The *SchemaSQL* Update Propagation Strategy

3.1 Classes of Updates and Transformations

The updates that can be propagated through *SchemaSQL* views can be grouped into two categories: *Schema Changes (SC)* and *Data Updates (DU)*. Schema changes that we consider are: *add-relation*(n, S), *delete-relation*(n), *rename-relation*(n, n') with relation names n, n' and schema S as introduced in Section 2.1 and *add-attribute*(r, a), *delete-attribute*(r, a), *rename-attribute*(r, a, a') with r the name of the relation R that the

attribute named a belongs to, a' the new attribute name in the rename-case, and the notation otherwise as above. Data updates are any changes affecting a tuple (and not the schema of the relation), i.e., *add-tuple(r,t), delete-tuple(r,t), update-tuple(r,t,t′)*), with t and t' tuples in relation R with name r. Note that we consider *update-tuple* as a basic update type, instead of breaking it down into a *delete-tuple* and an *add-tuple*. An *update-tuple* update consists of two tuples, one representing an existing tuple in R and the other representing the values of that tuple after the update. This is different from a sequence of *add* and *delete* as it allows to keep relational integrity constraints valid that would otherwise be violated temporarily.

3.2 *SchemaSQL* Update Propagation vs. Relational View Maintenance

Update propagation in *SchemaSQL*-views, as in any other view environment, consists in recording updates that occur in the input data, and translating them into updates to the view extent. In incremental view maintenance of SQL views [BLT86, QW91, GMS93, GL95, MKK97], many update propagation mechanisms have been proposed. Their common feature is that the new view extent is obtained by first computing *differences* between the old view V and the new view V' and then adding them to or subtracting them from the view, i.e., $V' = (V \backslash \nabla V) \cup \Delta V$, with ∇V denoting some set of tuples computed from the base relations that needs to be deleted from the view and ΔV some set that needs to be added to the view [QW91].

In *SchemaSQL*, this mechanism leads to difficulties. Even when considering only data updates to the base relations, executing the *SchemaSQL* view query V against the information space and update in the same way as above would lead to a ΔV or ∇V with unclear semantics. Since *SchemaSQL* views must propagate both schema *and* data updates, the schema of ΔV or ∇V does not necessarily agree with the schema of the output relation V. Also, the new view V' may have a different schema than V, which means that the concept of set difference between the tuples of V' and V is not even meaningful. An alternative formulation of the problem would be that any ΔV and/or ∇V computed in the above manner may not contain all the necessary information to update the view. At least, we must find a way to incorporate the concept of schema changes.

In order to overcome these difficulties, we now introduce a data structure ∂ which represents a sequence of n arbitrary data updates DU and schema changes SC.

Definition 2 (update sequence) *Assume two sets DU and SC which represent all possible data updates and schema changes, respectively, on a given relational schema R. Then an* update sequence *∂R is any arbitrary finite sequence of updates from DU and SC, i.e.,*

$$\partial R = (c_1, \ldots, c_n), \quad \forall i \in [1 \ldots n] : c_i \in DU \cup SC$$

Note that this definition extends naturally to views, since views can also be seen as relational schemas. For an example, consider propagation of the update `add-tuple('Berlin',1400,610) to` LH in Fig. 6 (p. 13). Having the value Berlin in the update tuple will lead to the addition of a new relation BERLIN in the output schema of the view—forming a sequence ∂V which contains both a schema change and a data update:

$$\partial V \quad = \quad (\textit{add-relation}(\mathsf{BERLIN}, (\mathsf{Type}, \mathsf{Destination}, \mathsf{BA}, \mathsf{LH})), \textit{add-tuple}(\mathsf{BERLIN}, (\text{'Economy'}, \mathsf{null}, 610)))$$

3.3 Overall Propagation Strategy

Given this new data structure ∂, and an implementation of it that we will call `Sequence`, our update propagation strategy works according to the algorithm in Fig. 4. Each node in the algebra tree has knowledge about the operator it represents. This operator is able to propagate *one* update and return a sequence ∂ of updates. Each operator can also recognize whether it is affected by an update (by comparing the relation(s) on which the update is defined with its own input relation(s)). If it is not affected, it simply returns an empty update sequence. We need this capability only for operators in leaf nodes, who must be able to determine whether an update occured in one of their base relations.

After all the updates for a node n's children are computed and collected in a single sequence s, they are propagated one-by-one through n itself. Each update entering n's operator will be propagated into a sequence of updates, all of which are concatenated into the final return sequence r.

```
function propagateUpdate(Node n, Update u)
    Sequence r ← ∅, s ← ∅
    if (n is leaf)
        if (n.operator is affected by u)
            r ← n.operator.operatorPropagate(u)
    else
        for(all children c_i of n)
            s ← s + propagateUpdate(c_i, u)
        for(all updates u_i in s)
            r ← r + n.operator.operatorPropagate(u_i)
    return r
```

Figure 4: The *SchemaSQL* View Maintenance Algorithm

The strategy to propagate an individual input change through the algebra tree to the final view corresponds to a postorder traversal of the algebra tree. In this way, each operator that needs to propagate updates works on input that is consistent with the current state of the information space, as all its children have already computed their output[2]. At each node N, an incoming update is translated into an output sequence ∂N of length $0 \ldots n$ which is then propagated to N's parent node. Since the algebra tree is connected and cycle-free (not considering joins of relations with themselves) all nodes will be visited exactly once. After all nodes have been visited, the output of the algorithm will be an update sequence ∂V to the view V that reflects the propagation of the incoming update.

3.4 Propagation of Updates through *SchemaSQL* Operators

Since update propagation in our algorithm occurs at each operator in the algebra tree, we have to design a propagation strategy for each type of operator. In this section, we will give expressions for the propagation of any update through any *SchemaSQL* operator.

A single incoming update can be propagated and transformed into one to many outgoing updates (cf. Figs. 1 and 7–10). As explained above, sequences ∂ of updates can be considered one update at a time. Therefore, we will consider update propagation only for single updates.

[2]We are not considering concurrent updates in this paper.

In Figs. 7–10, we give the update propagation tables for the four *SchemaSQL* operators. For the notation and meaning of variables and constants, please refer to Section 2.1. We also use the notation $\Rightarrow$ to denote the changing of the parameters of an operator. The tables show the propagation of each possible input change, broken down into several cases as necessary (in the second column). The second column also shows any auxiliary variables needed for propagation and their computation. The tables for UNFOLD and SPLIT contain a third column specifying affected objects, as the propagation of an update often depends on which attribute in a relation changes its values. For example, if an attribute is deleted from the input of a SPLIT-operator, the propagation depends on whether the deleted attribute is the pivot attribute A_p or another attribute. The last column in all tables shows the output of the operator, i.e., the update sequence $\partial\omega$ generated as the propagation of an input update. This output is shown as an SQL `insert`-, `delete`-, or `update`-statement (in `typewriter font`) or as a schema change (in *italics*). Note that whenever our propagation tables give an SQL-statement as the output of an operator ω, the exact list of updates $\partial\omega$ can be derived simply by applying this SQL-statement to the output relation R_ω. In the case of most SQL-`update` and `delete`-statements, the output relation of an operator ω is needed to translate a SQL-statement into an update sequence $\partial\omega$. This situation is similar to projection-operators in SQL, where the removal of duplicate output tuples can only be accomplished if the output of the operator is known. Otherwise, each operator ω in the tree requires as input for the propagation of any update only the input update itself, its own parameters, and the set A^* determined according to Section 2.2.

Inspection of the update propagation tables show several properties of our algorithm. For example, the view becomes invalid under some schema changes or data updates , mainly if an attribute or relation that was necessary to determine the output schema of the operator is deleted (e.g., when deleting the pivot or data attribute in UNFOLD). In other cases, an input schema change simply leads to the changing of a parameter in the corresponding operator. This is often the case if a schema element in the input relation is renamed (e.g., under *rename-relation* in FOLD). In this case, the operator will produce a zero-element output sequence $\partial\omega$.

3.4.1 Formalization of the Propagation of Updates

A formalization of the propagation of updates is extensive and lacks the conciseness of the propagation tables given in this section. Therefore, we will only give an example of how such a definition could be accomplished.

We will consider the propagation of *add-tuple* through UNFOLD:

Using the notation from Section 2.2.3, assume a relation $R = (N_R, S_R, E_R)$ with n attributes that is the input for an operator $Q = \text{UNFOLD}_{a_p,a_d}(R)$ producing an output relation $Q = (N_Q, S_Q, E_Q)$ with $n - 2 + k$ attributes and an update to R, denoted by $\Delta_R^+ = t[a_1, \ldots, a_n, a_p, a_d] = (x_1, \ldots x_n, x_p, x_d)$. Let A^* be a set of the k distinct values in A_p (defined above).

The propagation of this update is shown in Fig. 5. $(E_Q \setminus T_1)^*$ in the figure is constructed by *adding an attribute* to $E_Q \setminus T_1$, i.e., $(E_Q \setminus T_1 \subseteq D_1 \times \ldots \times D_{n+k}) \Rightarrow (E_Q^* \subseteq D_1 \times \ldots \times D_{n+k} \times D_d)$ with all data values in this new attribute set to NULL ($\perp$). Note that the output relation becomes invalid iff an update is inserted into the input relation that agrees in $a_1, \ldots, a_n, a_p$ with an existing tuple (similar to a key violation).

3.4.2 Propagation of Schema Changes through SQL Algebra Operators

The propagation of updates through standard SQL algebra nodes is simple. Deriving the update propagation for data updates is discussed in the literature on view maintenance [BLT86, QW91, GL95]. It remains to define update propagation for selection, projection, and cross-product operators under schema changes. As

$$
\begin{array}{rcl}
E'_R & \leftarrow & E_R \cup \Delta_R^+ \\
 & \Downarrow & \\
T_0 & \leftarrow & \{t \in E_R \mid t[a_1, \ldots, a_n, a_p] = \Delta_R^+[a_1, \ldots, a_n, a_p]\} \qquad \textit{find out if added tuple exists in } E_R \\
S'_Q & \leftarrow & \begin{cases} \emptyset & \text{if } T_0 \neq \emptyset \\ S_Q & \text{if } \Delta_R^+[a_p] \in A^* \\ (a_1, \ldots, a_n, a_1^*, \ldots, a_k^*, a_p) & \text{otherwise} \qquad \textit{schema change if necessary} \end{cases} \\
T_1 & \leftarrow & \{t \in E_R \mid t[a_1, \ldots, a_n] = \Delta_R^+[a_1, \ldots, a_n]\} \qquad \textit{tuples matching in key attributes} \\
T_2 & \leftarrow & \{t \in T_1 \mid \nu \leftarrow \Delta_R^+[a_p];\ t[\nu] \leftarrow \Delta_R^+[a_d]\} \qquad \textit{set pivot attribute to value in data attribute} \\
T_3 & \leftarrow & \begin{cases} \{t[a_1, \ldots, a_n, a_1^*, \ldots, a_k^*, a_p] \mid t[a_1, \ldots, a_n, a_p] \leftarrow \Delta_R^+[a_1, \ldots, a_n, a_d];\ t[a_i^*] \leftarrow \perp\} & \text{if } T_1 = \emptyset \\ \qquad \textit{if new row in output table, construct new tuple, fill unused attributes with NULL} & \\ \{t[a_1, \ldots, a_n, a_1^*, \ldots, a_k^*, a_p] \mid t[a_1, \ldots, a_n, a_1^*, \ldots, a_k^*] \in T_1;\ t[a_p] \leftarrow \Delta_R^+[a_d]\} & \text{otherwise} \\ & \textit{just set appropriate value} \end{cases} \\
E'_Q & \leftarrow & \begin{cases} \emptyset & \text{if } T_0 \neq \emptyset \\ (E_Q \setminus T_1) \cup T_2 & \text{if } \Delta_R^+[A_p] \in A^* \\ (E_Q \setminus T_1)^* \cup T_3 & \text{otherwise} \end{cases}
\end{array}
$$

Figure 5: Propagation of *add-attribute* through an UNFOLD-Operator

with all the other operators in this paper, we denote the input relation by R (and its name by r) and the output relation by Q (named q).

Relation-Updates: For all standard SQL algebra operators, *add-relation(n,S)* is not propagated since it will not affect existing output. *Delete-relation*(r) makes the operator's output invalid (i.e., is propagated into a *delete-relation*(q) of its output relation), and *rename-relation*(n, n') is not propagated since the name of an input relation is not relevant to a standard SQL-operator.

Attribute-Updates:

1. Selection operator σ_{cond} with *cond* a list of conditions on the attributes in this operator's relation R. All attribute updates are propagated by changing the parameter r (which represents the name of the input relation) to the name of the output relation q. If a predicate in *cond* is defined over a renamed attribute, it is changed accordingly. If it is defined over a deleted attribute, the view becomes invalid.

2. Projection operator $\pi_{\bar{A}}$ with a set of attributes $\bar{A}$ as projection list
   add-attribute(r, a) will not be propagated ($\partial\omega = \emptyset$) as a will not be in $\bar{A}$. *Delete-attribute*(r, a) will be propagated as *delete-attribute*(q, a) if $a \in \bar{A}$. *Rename-attribute*(r, a, a') will be propagated as *rename-attribute*(q, a, a') if $a \in \bar{A}$, in which case the projection list changes: $\bar{A}' \leftarrow \bar{A} \backslash a \cup a'$. Thus, the changed attribute name will be available on the output.

3. Cross-Product Operator $\times$
   The cross-product operator simply propagates any incoming attribute-update into an output update by changing the relation name parameter r to q.

3.5 Update Propagation Example

Fig. 6 gives an example for an update that is propagated through the *SchemaSQL*-algebra-tree in Fig. 2. The operators appear in boxes, with their output attached (SQL-statements according to our update tables in Figs. 7–10). The actual updates computed by these SQL-statements are shown in tabular form. The sending of updates to another operator is denoted by double arrows (⇑), while single arrows (↑) symbolize the transformation of SQL-statements that our update propagation tables suggest for each operator into updates.

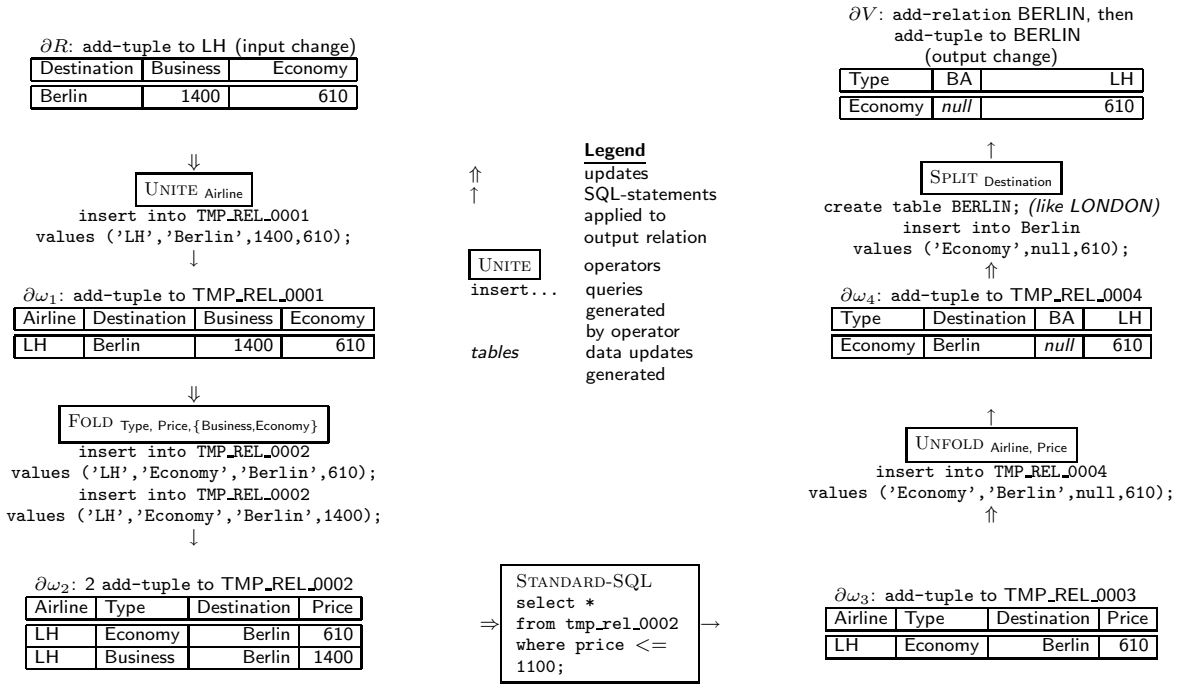

Figure 6: Update Propagation in the View from Figure 2. See Section 3.5 for explanation.

We are propagating an *add-tuple*-update in base relation LH through all five operators in the algebra tree. First, the UNITE operator propagates the incoming update into a one-element sequence $\partial\omega_1$ of updates which is then used as input to the FOLD-operator. The FOLD-operator propagates its input into a two-element sequence $\partial\omega_2$, sent to the StandardSQL-operator. This operator then propagates each of the two updates separately, creating two sequences $\partial\omega_{3_1}$ and $\partial\omega_{3_2}$, with 1 and 0 elements, respectively. Since one update is not propagated due to the WHERE-condition in the StandardSQL-node (i.e.,$|\partial\omega_{3_2}| = 0$), we need to show the further propagation of only $\partial\omega_{3_1}$. Recall from Section 3.3 that in the case of more than one update sequence being created by an operator, those sequences can simply be concatenated before the next operator's propagation is executed. UNFOLD now transforms its incoming one-element update sequence $\partial\omega_{3_1}$ into another one-element sequence $\partial\omega_4$ which becomes the input for the SPLIT-operator. This operator finally creates a two-element sequence, consisting of an *add-relation* schema change followed by an *add-tuple* data update. This sequence is the final update sequence ∂V which is applied to view V, leading to the new view V'. Note that one fact relevant for this particular propagation is that the value Berlin did not exist in any of the base tables before the update. Had the value existed, the propagation would not have included an *add-relation* schema change.

Input Change	Case	Object Affected	Propagation
add-tuple (r, t)	$t[a_1, \ldots, a_n, a_p] \in R$		invalid view
	$t[a_p] \in A^*$ $t[a_1, \ldots, a_n] \in R$		`update Q` `set` $[t[a_p]] = t[a_d]$ `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$
	$t[a_p] \in A^*$ $t[a_1, \ldots, a_n] \notin R$		`insert into Q` $(a_1, \ldots, a_n, a_p)$ `values` $(a_1, \ldots, a_n, a_d)$
	$t[a_p] \notin A^*$		*add-attribute(q,*$t[a_p]$*)*, `update Q set` $[t[a_p]] = t[a_d]$ `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$
delete-tuple (r, t)	$t[a_p]$ exists in A^* exactly once		*delete-attribute(q,*$t[a_p]$*)*
	$t[a_p]$ exists in $A^* >$ once		`update Q set` $[t[a_p]] =$ `NULL` `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$ [3]
update-tuple (r, t, t')		A_d	`update Q set` $[t[a_p]] = t[a_d]$ `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$
		any other att.	break down into (*delete-tuple*, *add-tuple(*
add-attribute(r, a)			*add-attribute*(q, a)
delete-attribute(r, a)		A_d or A_p	invalid view
		other	*delete-attribute*(q, a)
rename-attribute(r, a, a')		A_d	$\textsc{Unfold}_{a_p,a} \Longrightarrow \textsc{Unfold}_{a_p,a'}$
		A_p	$\textsc{Unfold}_{a,a_d}(R) \Longrightarrow \textsc{Unfold}_{a',a_d}(R)$
		other	*rename-attribute*(q, a, a')
delete-relation(r)			*delete-relation*(q)
rename-relation(n, n')			$\textsc{Unfold}_{a_p,a_d}(N) \Longrightarrow \textsc{Unfold}_{a_p,a_d}(N')$

[3] if this update leads to a tuple with all NULL-values, the tuple must be deleted.

Figure 7: Propagation Rules for Q=$\textsc{Unfold}_{a_p,a_d}(R)$

Input Change	Case and Variable Binding	Propagation
add-tuple (r, t)	$(A^* = \{a_1^*, \ldots, a_k^*\}, k \leftarrow \lvert A^* \rvert)$	`for` i `:= 1..`k `insert into Q` `values` $(a_1, \ldots, a_n,$`,`$a_i^*,$`,`$t[a_i^*])$
delete-tuple (r, t)		`delete from Q` `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$
update-tuple (r, t, t')	$A \in A^*$; set $t[a]$ to a value c	`update Q set` $a_d = c$ `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$ `and` $a_p = a$
	$A \notin A^*$; set $t[a]$ from a value b to a value c	`update Q set` $a = c$ `where` $a = b$
add-attribute(r, a)	$A \in A^*$ [4]	`foreach tuple` $u \in R$ `insert into Q` $(a_1, \ldots, a_n, a_p, a_d)$ `values` $(u[a_1, \ldots, a_n], a,$`NULL`$)$
	$A \notin A^*$	*add-attribute*(q, a)
delete-attribute(r, a)	$A \in A^*$	`delete from Q where` $a_p = a$
	$A \notin A^*$	*delete-attribute*(q, a)
rename-attribute(r, a, a')	$A \in A^*$	`update Q set` $a_p = a'$ `where` $a_p = a$
	$A \notin A^*$	*rename-attribute*(q, a, a')
delete-relation(r)		*delete-relation*(q)
rename-relation(n, n')		$\textsc{Fold}_{a_p,a_d}(N) \Longrightarrow \textsc{Fold}_{a_p,a_d}(N')$

[4] Note that the decision whether a *new* attribute should be a member of $a_1, \ldots, a_n$ can only be made by evaluating the view query.

Figure 8: Propagation Rules for Q=$\textsc{Fold}_{a_p,a_d,A^*}(R)$

Input Change	Case	Object Affected	Propagation
add-tuple (r, t)	$t[a_p] \notin A^*$		`add-relation` $[t[a_p]]$ `with schema` $(S_R \backslash R.A_p)$; `insert into` $[t[a_p]]$ `values` $(t[a_1, \ldots, a_n])$
	$t[a_p] \in A^*$		`insert into` $[t[a_p]]$ `values` $(t[a_1, \ldots, a_n])$
delete-tuple (r, t)	$t[a_p]$ exists in A^* exactly once		`delete-relation` $[t[a_p]]$
	$t[a_p]$ exists in $A^* >$ once		`delete from` $[t[a_p]]$ `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$
update-tuple (r, t, t')		A_d	`update` $[t[a_p]]$ `set` $[a_d]$ `=` $t[a_d]$ `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$
		any other att.	break down into (*delete-tuple*,*add-tuple*)
add-attribute(r, a)			$\forall q \in \{q_1 \ldots q_n\}$: *add-attribute*(q, a)
delete-attribute(r, a)		A_p	invalid view
		other	$\forall q \in \{q_1 \ldots q_n\}$: *delete-attribute*(q, a)[5]
rename-attribute(r, a, a')		A_p	$\textsc{Split}_a(R) \Longrightarrow \textsc{Split}_{a'}(R)$
		other	$\forall q \in \{q_1 \ldots q_n\}$: *rename-attribute*(q, a, a')
delete-relation(r)			$\forall q \in \{q_1 \ldots q_n\}$: *delete-relation*(q)
rename-relation(n, n')			$\textsc{Split}_{a_p}(N) \Longrightarrow \textsc{Split}_{a_p}(N')$

[5]If this update leads to a tuple with all NULL-values in an output relation, the tuple must be deleted.

Figure 9: Propagation Rules for Q=$\textsc{Split}_{a_p}(R)$

Input Change	Case and Variable Bindings	Propagation
add-tuple (r_x, t)		`insert into Q` `values` $(t[a_1, \ldots, a_n], r_x)$
delete-tuple (r_x, t)		`delete from Q` `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$ `and` $a_p = r_x$
update-tuple (r_x, t, t')	$A = A_d$; set $t[a]$ to a value c	`update Q set` $a = c$ `where` $a_1, \ldots, a_n = t[a_1, \ldots, a_n]$ `and` $a_p = r_x$
	$A \neq A_d$; set $t[a]$ from a value b to a value c	`update Q set` $a = c$ `where` $a = b$ `and` $a_p = r_x$
add-attribute(r, a)	add simultaneously to all R_i	*add-attribute*(q, a)
	otherwise	invalid view
delete-attribute(r, a)	delete simultaneously from all R_i	*delete-attribute*(q, a)
	otherwise	invalid view
rename-attribute(r, a, a')	rename simultaneously in all R_i	*rename-attribute*(q, a, a')
	otherwise	invalid view
add-relation(r_x, S)		no change (until first *add-tuple* to R_x)
delete-relation(r_x)		`delete from Q where` $a_p = r_x$
rename-relation(n, n')		$\textsc{Unite}_{a_p}(\{R_1, \ldots, N, \ldots, R_n\}) \Longrightarrow \textsc{Unite}_{a_p}(\{R_1, \ldots, N', \ldots, R_n\})$

Figure 10: Propagation Rules for Q=$\textsc{Unite}_{a_p}(R_1, R_2, \ldots, R_n)$[6]

[6]Note that r_x is the name of Relation R_x, which is one of the n relations of equal schema that are united by the Unite-operator.

3.6 Grouping Similar *SchemaSQL* Updates in Batches

Certain updates in our strategy are transformed by some operators into update sequences ∂ in which all the updates are similar. This gives an opportunity for optimization on our update propagation strategy.

For example, a FOLD-node can transform a single schema change (such as a *attribute-delete*) into a sequence of data updates (such as a sequence of *tuple-deletes*). An inspection of the update propagation tables in this section shows that, typically, such a sequence consists of similar updates. Consider the example in Fig. 6, where a deletion of attribute Business in relation TMP_REL_0001 would lead to a sequence of *delete-tuple* updates of all tuples in TMP_REL_0002 that have the value Business in the attribute Type. A simple way of executing all those updates efficiently using SQL would be to issue a query such as `delete from TMP_REL_0002 where type='Business'`. Thus, instead of propagating all individual tuple updates using some delta relation, as done in traditional view maintenance, we instead propose to abstract this sequence of updates into an SQL update statement and push the complete statement through the algebra tree.

We have identified two classes of such *batched updates* that occur frequently as outputs of our propagation strategy, as described below.

Definition 3 (Batched Update) *A* batched update *is a sequence ∂ of* SchemaSQL *updates which adheres to one of the following structures:*

- *∂ consists entirely of* delete-tuple-*updates to the same relation R, with equal schema and a set of attributes $a_1 \dots a_k$ whose values are a unique identifier for each tuple in ∂. We denote such a sequence by* delete-tuple-batch$(r, cond(a_1, c_1), \dots, cond(a_k, c_k))$ *with $cond(a, c)$ a condition selecting tuples $t \in R$ that have value c in attribute a ($t[a] = c$). This represents a set of delete-tuple statements on the output relation R that could be generated by an SQL-`delete` statement with the WHERE-conditions* $\mathsf{a_1} = \mathsf{c_1}, \dots, \mathsf{a_k} = \mathsf{c_k}$.

- *∂ consists entirely of* update-tuple-*updates to relation R, with equal schema. Parameters are a single attribute b with (unique or duplicate) values, and a function f between the old and new values of another attribute a in each tuple. We denote such a sequence by* update-tuple-batch(r, a, f, b, c) *with a, b attribute names, f a function over the domain of a ($f : D_a \rightarrow D_a$), and c a constant. This represents a set of update-tuple updates affecting every tuple t for which the value of attribute b is c, by changing the value $t[a]$ to $f(t[a])$, i.e., $\forall t \in R$ s.t. $t[b] = c : \ t[a] \leftarrow f(t[a])$.*

We do not define insert-tuple batches since we consider only single data updates or schema changes entering our algebra tree, and such updates will not be transformed into larger "batches". In particular, adding an attribute or a relation in a base table means adding an *empty* structure, so that such updates do not lead to batches of updates, and in fact often do not lead to any updates on the view extent at all. Batches of schema-changes are also not useful because meaningful schema-change batches do not occur in our context. Inspection of the update tables in this section shows that, with the exception of the SPLIT node, propagation of schema changes always leads to a *single* schema change, not sequences of related changes. In the case of the SPLIT-node, any resulting "batch" of schema changes will lead to changes across several relations, an operation that cannot be optimized using SQL-statements.

With this definition of batched update, the above example can now be represented as *delete-tuple-batch* (TMP_REL_0002, *cond*(type, 'Business')).

As mentioned above, the main benefit of batched updates lies in a possible optimization of the implementation of our update propagation strategy. Since some operators generate batches of related updates,

considering batches as types of updates and propagating those through the algebra tree just like single updates could lead to performance improvements of the system. For an example, consider again Fig. 6 and an update *delete-attribute*(TMP_REL_0001, Business) as input to the FOLD-operator. Setting $n = |\text{TMP_REL_0001}|$, this update in the current strategy would lead to a propagation of n single *delete-tuple* updates, whereas a treatment of all those updates as a batch would require the propagation of only one update, namely *delete-tuple-batch*(TMP_REL_0002,*cond*(type,'Business')). Figures 11–14 show the propagation tables.

Input Change	Parameters	Addtl. Conditions	Propagation
delete-tuple-batch	$(r, cond(a, c))$	$a = a_p$	`del-att` $[c]$

		$a = a_d$, A^* unchanged	`foreach` $a \in A^*$
			`update Q`
			`set` $a =$ `NULL`
			`where` $a = c$

		other	*delete-tuple-batch* $(q, cond(a, c))$
	$(q, cond(a_1, c_1), \ldots, cond(a_n, c_n))$		*delete-tuple-batch* $(q, cond(a_1, c_1), \ldots, cond(a_n, c_n))$
update-tuple-batch	$(r, a_p, f, a_p, c_{\text{old}})$ with $f(v) = c_{\text{new}}$ (a constant function)		`rename-att` c_{old} `to` c_{new}
	(r, a_d, f, b, c)		`foreach` $a \in A^*$ `update Q` `set` $[a] = f([a])$ `where` $b = c$

Figure 11: Batched Update Propagation Rules for Q=UNFOLD $_{a_p,a_d}(R)$

Input Change	Parameters	Addtl. Conditions	Propagation
delete-tuple-batch	$(r, cond(a, c))$	$a \in A^*$	`delete from Q` `where` $a_p = a$ `and` $a_d = c$
		other	*delete-tuple-batch* $(q, cond(a, c))$
	$(q, cond(a_1, c_1), \ldots, cond(a_n, c_n))$		*delete-tuple-batch* $(q, cond(a_1, c_1), \ldots, cond(a_n, c_n))$
update-tuple-batch	(r, a, f, b, c)	$a \in A^*$	`update Q` `set` $a_d = f(a_d)$ `where` $a_p = a$ `and` $b = c$

Figure 12: Batched Update Propagation Rules for Q=FOLD $_{a_p,a_d,A^*}(R)$

4 Correctness

Our update propagation strategy is equivalent to a stepwise evaluation of the algebraic expression constructed for a query. Each operator transforms its input changes into a set of semantically equivalent output changes, eventually leading to a set of changes that must be applied to the view to synchronize it with the base relation change. In this section, we will show that this strategy leads to correct update propagation. Recall that we denote the output update sequence of operator ω by $\partial\omega$ (Definition 2). We will use the notation R_ω for the input relation and Q_ω for the output relation throughout this section.

Theorem 1 (Correctness of Incremental Update Propagation for *SchemaSQL* Operators) *Let* $\omega \in \{\text{UNITE}, \text{SPLIT}, \text{UNFOLD}, \text{FOLD}, \pi, \sigma, \times\}$ *be a node in a* SchemaSQL *algebra tree. Let* R_ω *be the input relation(s) for* ω *and* $Q_\omega = \omega(R_\omega)$ *be its output relation(s). Furthermore, let* ΔR_ω *be a data update or schema change to* R_ω*, transforming* $R_\omega \rightarrow R'_\omega$*. Then, applying the rules from the update propagation tables*

Input Change	Parameters	Addtl. Conditions	Propagation
delete-tuple-batch	$(r, cond(a, c))$	$a = a_p$	`del-rel` $[c]$
		$a = a_d$, A^* unchanged	`foreach` $q \in A^*$ `update` q `set` $q.a_d =$ NULL `where` $q.a_d = c$
		other	`foreach` $q \in A^*$ *delete-tuple-batch*$(q, cond(a, c))$
	$(q, cond(a_1, c_1), \ldots, cond(a_n, c_n))$		`foreach` $q \in A^*$ *delete-tuple-batch* $(q, cond(a_1, c_1), \ldots, cond(a_n, c_n))$
update-tuple-batch	$(r, a_p, f, a_p, c_{\text{old}})$ with $f(v) = c_{\text{new}}$ (a constant function)		`rename-rel` c_{old} `to` c_{new}
	(r, a, f, b, c)	$a \in A^*$	`foreach` $q \in A^*$ *update-tuple-batch*(q, a, f, b, c)

Figure 13: Batched Update Propagation Rules for Q=SPLIT $_{a_p}(R)$. Note that $A^* = \{q_1, q_2, \ldots, q_k\}$ is the set of output relation names.

Input Change	Parameters	Addtl. Conditions	Propagation
delete-tuple-batch	$(r_x, cond(a, c))$		`delete from Q` `where` $a_p = r_x$ `and` $a = c$
update-tuple-batch	(r_x, a, f, b, c)	$a \in A^*$	`update Q` `set` $a = f(a)$ `where` $a_p = r_x$ `and` $b = c$

Figure 14: Batched Update Propagation Rules for Q=UNITE $_{a_p}(R_1, R_2, \ldots, R_n)$

Figs. 7–10 and Section 3.4.2 for ω and ΔR_ω will transform ΔR_ω into a sequence of updates to the node's output relation (denoted by ∂Q_ω) that transforms $Q_\omega \rightarrow Q'_\omega$.

Proof: For $\omega \in \{$UNITE, SPLIT, UNFOLD, FOLD$\}$ the proof is by inspection of the update propagation tables, Figs. 7–10. For the remaining operators, the propagation rules in Section 3.4.2 can be easily verified. □

Theorem 2 (Closure of Update Propagation) *Let DU and SC be the set of all possible data updates and schema changes, respectively, as defined in Section 3.2. Then, for each SchemaSQL or SQL operator $\omega \in \{$UNITE, SPLIT, UNFOLD, FOLD, $\pi, \sigma, \times\}$, a total function $f_\omega : (DU|SC) \rightarrow (DU|SC)^*$ can be defined that generates a correct output sequence ∂. That is, for every operator used in our algebra, the algorithm defined by our update propagation tables for each operator takes any data update or schema change and produces a sequence of zero to many changes, denoted by ∂, that consists entirely of data updates from DU and/or schema changes from SC.*

Proof: The proof is by inspection of the update propagation tables and the arguments in Section 3.4.2. □

Note that the above theorem means that for each operator, our update propagation tables define functions *closed under update sequences* ∂. It follows that update propagation can occur through the entire algebra tree, not only through a single operator, since the output of each operator's update propagation can be "fed" into another operator's propagation algorithm as input. This notion is formalized in the following theorem:

Theorem 3 (Closure under Update Sequences) *Let ω_1 be a* SchemaSQL *algebra operator defined over R, ω_2 an operator defined over $\omega_1(R)$, and $Q = \omega_2(\omega_1(R))$. That is, the schema of R and the output schema of ω_1 comply with the expected input schemas of ω_1 and ω_2, respectively. Then, the following holds:*

(1) The change sequence ∂_{ω_1} generated by ω_1 under any input schema change ΔR (transforming R to R') can be used as an input to ω_2.

(2) In any algebra tree for the evaluation of a SchemaSQL*-view V, any single update (data update or schema change) to a base relation can be propagated through the entire tree to finally produce an update sequence ∂V to the view V.*

Proof: The proof for the first part of the theorem follows immediately from Theorem 2 since any operator produces only output from $\{DU, SC\}$ which any other operator can process. The second part follows by repeated application of the first part. □

Theorem 4 (Correctness of *SchemaSQL* View Maintenance) *Let V be a view defined over the set of base relations $R_1, \ldots, R_n$, and $\Delta R_i \in \{DU, SC\}$ an input update applied to one relation R_i $(1 \leq i \leq n)$. Let R_i' be the relation R_i after the application of ΔR_i. Furthermore, let the* SchemaSQL *View Maintenance Algorithm as defined in Section 3.3 produce a change sequence ∂V that transforms view V into view V'. Then, view V' is equivalent in schema and extent to the view V^* obtained by recomputation over base relations $R_1, \ldots, R_i', \ldots, R_n$*

Proof: The proof is by recursive simplification of the algebra expression corresponding to the algebra tree. It is to prove that the update propagation through the innermost operator in the algebra expression (Equation 1) leads to a correct output relation. Note that this operator is a leaf node in the algebra tree. We also have to prove that the sequence of propagated updates from the innermost operator will correctly update the output relation of the next operator in the algebra tree, continuing this process through the entire tree until the root (i.e., the view) is reached.
According to Equation 1 (Section 2.2.6), a *SchemaSQL* algebra expression has the following general form (with indexed algebra nodes ω) :

$$V = \omega_1(\omega_2(\ldots\omega_j[\omega_{k-1}(\omega_k(R_1)), \ldots, \omega_{n-1}(\omega_n(R_m))]\ldots)) \tag{3}$$

with $\omega_i \in \{\textsc{Unite}, \textsc{Split}, \textsc{Unfold}, \textsc{Fold}, \pi, \sigma, \times\} \cup \mathcal{I}$ and $\mathcal{I}$ the identity function (i.e., any operator can be omitted for a particular query). Note that ω_j corresponds to the $\times$-operator in Equation 1. Lakshmanan et al. [LSS99] show that this expression can be used to correctly evalute a *SchemaSQL* query. Now each operator ω_i computes an output relation Q_{ω_i} (if ω is a Split-operator, it may produce several relations of the same schema). In particular, ω_n computes a relation based only on the data in R_m. ω_{n-1} computes a relation based only on the data in Q_{ω_n}, the output relation of operator ω_n and so on. We will define the relations computed by each operator ω_i to be *correct* for *SchemaSQL* query evaluation since we assume that the initial query evaluation (as well as all recomputations) produce the correct overall result for the view and must therefore also produce correct results at each *node* in the algebra tree.

Note also that the structure of the algebra tree for a view depends only on the query, not on the base data [LSS99]. The only changes to operators under base relation updates are possible changes of parameters inside the operators; an algebra operator can not disappear or appear as the result of a base update [7].

Now consider an update ΔR_i to a base relation R_i, transforming R_i into its new state R_i' as defined above. If we were to recompute the view over the new state of the information space, the view extent V' obtained by recomputation would be

$$V' = \omega_1(\omega_2(\ldots\omega_j[\omega_{k-1}(\omega_k(R_1)), \ldots, \omega_{p-1}(\omega_p(R_i')), \ldots, \omega_{n-1}(\omega_n(R_m))]\ldots)) \tag{4}$$

[7] However, the view may become invalid under some updates.

After this recomputation, each operator ω_i will have produced an output relation Q'_{ω_i} that is defined as correct (since we define recomputation to produce the correct result of the query evaluation). Since we want to prove equivalence of our propagation result to recomputation, we need to show that all output relations produced by the *SchemaSQL* View Maintenance strategy are equivalent to the relations Q'_{ω_i} produced by algebra tree recomputation. This will prove our theorem since $V' \equiv Q'_{\omega_1}$.

(1) With Theorem 1, we can conclude that operator ω_p produces a sequence of changes ∂Q_{ω_p} to its output relation Q_{ω_p} that transforms it into the correct relation Q'_{ω_p}.

(2) Note that during view materialization, operator ω_{p-1} has correctly computed its output relation $Q_{\omega_{p-1}}$ from its input $R_{\omega_{p-1}}$. Due to Theorem 1, ω_{p-1} will, in an update propagation step, correctly transform its output relation $Q_{\omega_{p-1}}$ to the new state $Q'_{\omega_{p-1}} = \omega_{p-1}(R'_{\omega_{p-1}})$ when given the correct sequence of input changes $\partial R_{\omega_{p-1}}$ that transforms $R_{\omega_{p-1}} \rightarrow R'_{\omega_{p-1}}$. But, since $R_{\omega_{p-1}} \equiv Q_{\omega_p}$, this sequence is equal to ∂Q_{ω_p} which has been established in step (1). Thus, it is possible to apply it to ω_{p-1} (due to Theorem 3) and it will lead to the expected output.

(3) Steps (1) and (2) are executed subsequently for all operators $\omega_{p-1}, \ldots, \omega_1$ whose input update sequence is not empty. That is, each operator ω_i will transform its (correct) output relation Q_{ω_i} according to the changes to its input relation R_{ω_i}) into the correct output relation Q'_{ω_i}, such that eventually operator ω_1 will produce a sequence $\partial Q_{\omega_1} \equiv \partial V$ that correctly transforms the view V to the new view V'.

Since an update occurs in exactly one base relation, and the algebra tree has no cycles, no operator will receive input update sequences from more than one of its children, avoiding concurrency issues. □

5 Implementation

5.1 *SchemaSQL* Query Engine

The update propagation strategy described in this paper has been implemented in Java on top of a *SchemaSQL* query evaluation module also written by us. This query engine was built along the lines of [LSS99]. Our code first parses *SchemaSQL* queries (using JavaCC), then builds an algebra tree out of the parsed query, and finally evaluates the query result through a postorder traversal of that tree, computing the output of each algebra node as it is visited (cf. Fig. 15). The next node then reads its child node's temporary relation to compute its output. For this prototype implementation, each node temporarily stores its output through JDBC in the query engine's "local" relational database (Oracle 8) as keeping relations in memory only incurs limitations on the size of input relations and also would have required us to reimplement significant parts of relational query technology. For performance reasons, all intermediate nodes in the algebra tree share one JDBC-connection to the local database.

The implementation of the query engine uses pure Java and JDBC-connections to several instances of Oracle 8. We use standard SQL DDL and DML statements (`select`, `insert`, `delete`, `update` and statements for schema change operations like `alter table`) for all queries—thus making full use of the source database's SQL query evaluation capabilities. We do not use any system specific functions other than simple schema changes. A wrapper class (which we have also successfully implemented for Microsoft's Access) makes differences in the syntax of schema change operations transparent. Therefore, the implementation is independent of the database used.

To improve performance and simplify our code, our implementation attempts to use as much of the SQL query capability as possible, in particular by extracting standard SQL out of the *SchemaSQL* query

and evaluating it in a single `StandardSQL`-node, which simply executes its stored SQL-statement against the local SQL-database. The *SchemaSQL* query engine currently does not perform or utilize any query optimization strategies other than those provided by the underlying SQL query engine when executing queries against the local database.

5.2 Incremental Update Propagation

To perform update propagation in the manner described in this paper, we added update propagation capabilities to each algebra operator class (UNITE, FOLD, UNFOLD, SPLIT, StandardSQL). A method `propagateUpdate()` in each node accepts one update and returns a list of (data and/or schema) updates which represent the result of the update propagation. Then, we added code for the propagation of the output updates of each operator to its parent. Thus, the same code that performs the postorder traversal of the operator tree for the initial materialization of the view can now also perform the incremental update propagation, by simply calling the update propagation method instead of the materialization method on each node and recursively using each operator's output as the input for the parent operator.

Updates were modeled into a small class hierarchy consisting mainly of the classes `SchemaUpdate` and `DataUpdate`. Each node in the algebra tree is now extended by the ability to propagate all such updates. The current code does not support batched updates.

Fig. 15 shows the architecture of our system.

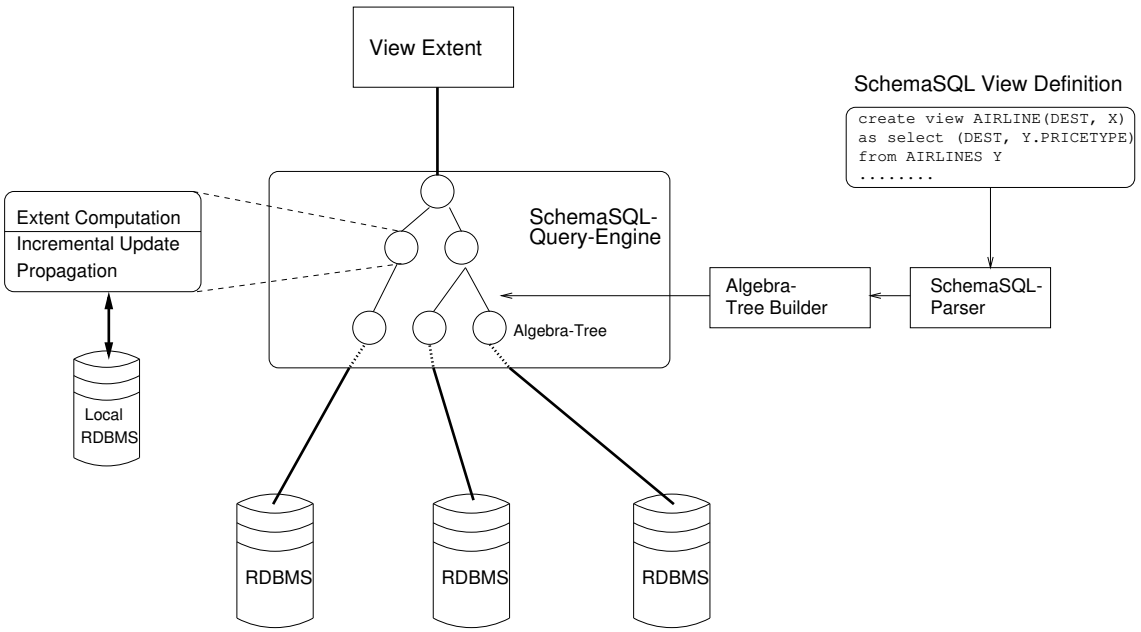

Figure 15: The architecture of the *SchemaSQL* View Maintenance System

6 Performance Evaluation

6.1 Experimental Setup

6.1.1 Performance-Relevant Factors

A explained in this paper, *SchemaSQL* update propagation is significantly different from traditional incremental view maintenance. Major differences are the transformation of data updates into schema changes and vice-versa, the need for the propagation of base schema changes, and the propagation strategy based on propagating an update through an algebra tree rather than computing delta-queries against the base relations.

To assess the influence of those issues on performance, we executed a number of experiments on our prototype. We are reporting some of the results in this section. For the experiments described here, we focused on the following factors contributing to *SchemaSQL* update propagation performance:

- the type of update (data update or schema change) at the base relations;
- the transformation type of the update (i.e., the type, data update or schema change, into which a base update is propagated);
- the selectivity of conditions in the view query that determines the size of the view relative to the sizes of the base relations;
- the size of base relations.

6.1.2 Schema and View Queries

If not stated otherwise, all our experiments use the following view query, over the same base schema as in our running example (Fig. 1):

```
create view CITY(Type, AIRLINE) AS
  select PRICETYPE, FLIGHT.PRICETYPE
  from   -> AIRLINE,
         AIRLINE FLIGHT,
         AIRLINE-> PRICETYPE,
         FLIGHT.Destination CITY
  where PRICETYPE <> 'Destination'
  and AIRLINE like 'AIRLINE%'
  and FLIGHT.PRICETYPE <= '1101';
```

The output schema of this query, considering the input schemas from Fig. 1, consists of two relations Business and Economy which both have the schema (Destination,AIRLINE_1,...,AIRLINE_K), with one attribute named AIRLINE_X for each relation named AIRLINE_X in the input schema. The output schema may change during an experiment.

The base data was generated from a list of strings (representing city names), augmented by random number representing "flight prices". Those numbers were generated using as uniformly distributed random numbers in a certain range. The base relation sizes and distribution of updates are described with each experiment.

Since we have multiple output relations, we need to extend the concept of view size to multiple relations. We thus define the *view size* to be the *sum* of the sizes of all output relations.

6.1.3 Measurements and System Parameters

The test system was a Pentium II/400 running Linux and Java 1.2.2. Database connectivity was achieved through JDBC. The databases used for our tests (local database and information source) were two installations of Oracle 8i, running on a Pentium 233 under Windows NT and a 4-processor 300MHz DEC Alpha under DEC OSF1, respectively.

For our experiments, we measured the total execution time of the initial materialization of the view (for control purposes, not shown in the chart), then the time for a number of updates, depending on the experiment, and finally the time for a recomputation of the view extent. The times were measured by comparing the system time before and after executing update propagation or recomputation, i.e., they include system, user, and I/O time.

6.2 Deleting Base Relations of Different Sizes

In our classification of updates (Section 3), inserting a table implies inserting an *empty* table. The data would have to be added in subsequent data updates. Therefore, inserting (and also renaming) schema elements leads to relatively simple propagation results, as the information content of the database does not change much under such updates. Therefore the experiments reported here concentrate mainly on the deletion of schema elements (attributes, relations), as well as on the insertion and deletion of data.

We ran the above query over a schema containing four base relations (representing four different airlines) with approximately 111, 222, 333, and 444 tuples, respectively. We then deleted each of those base relations and compared the time for incremental update propagation with the time for recomputation after a base relation was deleted. After each deletion and measurement, the original information space was restored. With the above query, the original view extent had two relations Business and Economy with 378 and 444 tuples, respectively and decreased roughly proportionally after deletions of base relations. Fig. 16 shows the times measured.

In our current implementation, deleting a base relation R in a query like above will result in the creation of approximately $|R|$ updates inside the operator tree. Therefore, deleting larger relations takes longer than deleting smaller relations. On the other hand, recomputation time will decrease with larger sizes of the deleted relation, as the resulting view extent has less tuples. The crossover point between view maintenance and recomputation is at about a base size of 225, i.e., deleting a relation of more than 225 tuples (about 25% of the view size) will take longer than recomputing the view. This means that our update propagation strategy will perform better than recomputation for tuple-wise deletions of up to 25% of the entire information space. Propagating 225 delete-tuple updates will have a smaller total execution time than a single recomputation of the view. On the other hand, this experiment shows that a single input update (in this case a *delete-relation* update) can be very expensive as it may lead to many updates in the view extent.

We also ran the same set of updates under the view query from our running example (Fig. 1), which creates one output relation for each unique destination in any of the input relations. A first experiment showed that, during the materialization of the view from large base relations, the Oracle database we were using rejected the creation of new tables after about 620 `create table`-statements. So we repeated the experiment with the above schema having 111,222,333, and 444 tuples and deleted each relation. This time, the query computed properly. Some of the running times are given in Table 1.

Note that in this case, the deletion of a base relation eventually leads to *one delete-attribute schema change per output relation*, so deleting a base relation with 222 tuples will first lead to the propagation of

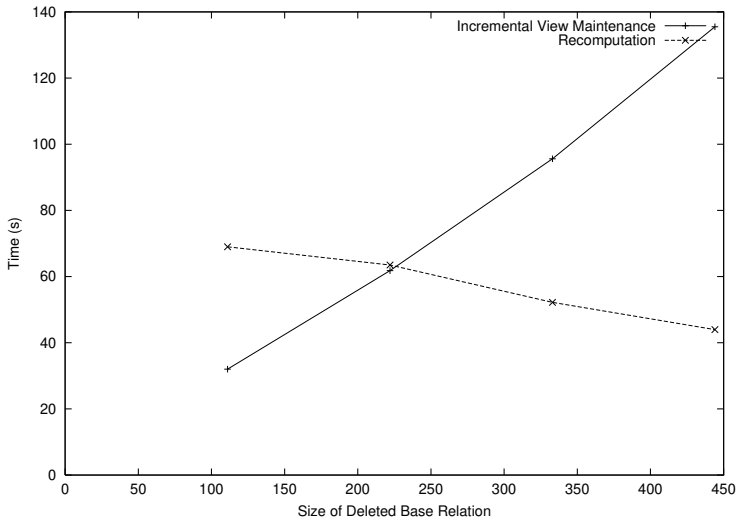

Figure 16: View Maintenance and Recomputation Times vs. Size of Deleted Base Relations

Base Table Deleted	Relation Size	Incremental VM	Recomputation
AIRLINE_LH	333	395s	125s
AIRLINE_KL	111	318s	145s

Table 1: Propagation and Recomputation in the Query from Fig. 1

222 *delete-tuple* updates, propagated to one output relation each, and then to up to 444 *delete-attribute* schema changes—exactly one for each output relation. The experiment shows that in this case, which incurs many schema changes in the output relations, a recomputation has a performance advantage over the unoptimized incremental update propagation. An obvious optimization for the processing of this particular query would be to ignore the *delete-tuple*-updates and apply only the *delete-attribute* changes.

6.3 Deleting Tuples from Base Relations

In this experiment (Fig. 17), we delete a sequence of random tuples from the base schema and measured the cumulative propagation time. For our chart, we numbered the updates by consecutive numbers i. The cumulative propagation time for update i is the sum of the propagation times for all updates numbered $0 \ldots i$. We also measured the time to recompute the view after the entire update sequence was executed. This experiment shows again that in our schema, the crossover point between incremental maintenance and recomputation is at roughly 200 tuples, i.e., after 200 updates, recomputation of our view (with view size 888) would become more efficient. The view size is the major factor determining the recomputation time for a view, whereas the time for incremental propagation of a single data-update mainly depends on the system implementation, i.e., is roughly constant for our implementation and test environment. The average time to propagate a single update can be estimated from the slope of the curve in Fig. 17 to be about 285ms, which is a value depending mainly on the size of the tuples propagated.

From those facts, it can be concluded that the *ratio* between the number of propagated updates at the crossover point and the view size is a system constant depending only on the implementation, i.e., we expect that recomputation of a view will take roughly the same time as the incremental propagation of the deletion

of a certain percentage of the view's tuples. In our experiment, this ratio was about 1/4 (a crossover point of 200 for a view with roughly 800 tuples).

Note the jump in the incremental maintenance time at the end of the curve. This time represents a data update that led to a schema change in the output relation. The reason is that the last tuple from a base relation was deleted which led to a *delete-attribute* change in an output relation.

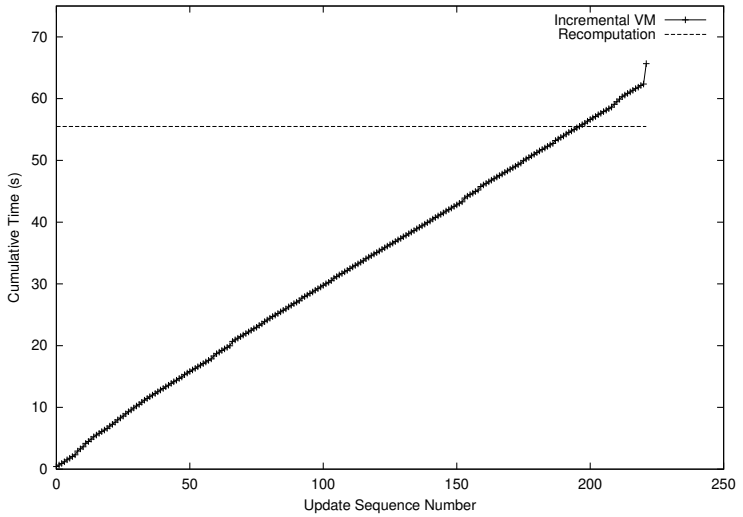

Figure 17: Deleting Tuples from the Input Schema

6.4 Deleting Tuples Leading to Schema Changes

In this experiment, we wanted to assess the difference in propagation time of the same updates, depending on whether these updates lead to data updates or schema changes in the view. Schema changes actually executed on a relation database are slow operations. Therefore, the expectation is that an update propagation (including the application of those updates against a database) that leads to a schema change will be slower than an update propagation leading to only a data update in the view.

Thus, in this experiment we are deleting tuples from the base schema. This time we have four base relations R_1, R_2, R_3, R_4 of sizes 1, 10, 100, and 1000, respectively, but make sure that some of the updates incur schema changes in the output schema. First, we inserted 10 tuples into relation R_1, then removed tuples from relation R_2 one-by-one, until R_2 was empty. This leads to a schema change in the output schema since the corresponding attribute is removed from each output relation. Then, we removed all 11 tuples from relation R_1 (incurring another schema change), followed by a removal of 10 random tuples from relation R_3. Note that in this sequence, updates #19 and #30 lead to schema changes in the output.

We then plotted the time each update took to propagate. Note the relatively even distribution of update propagation times around 250 ms in Fig. 18, except for two updates, which take over 2 seconds to propagate. Those are clearly those updates that led to schema changes in the output relations. The measured update time (if no schema change is incurred) is similar to the average time measured in Fig. 17. Again, the propagation time for the data updates depends mainly on the tuple size, and the underlying database, whereas the time for schema changes depends on the underlying database only (as practically no data has to be transported).

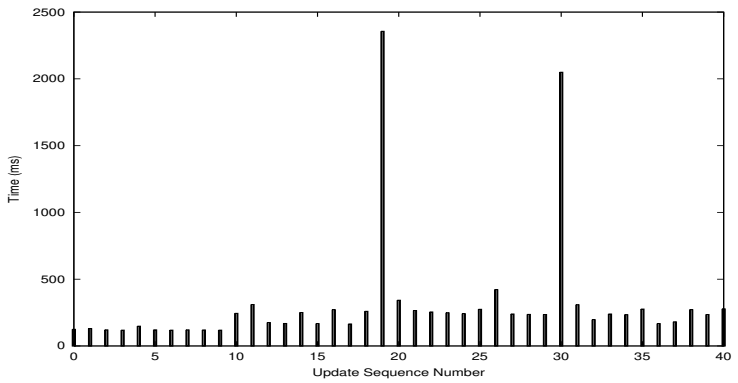

Figure 18: Base Updates lead to Data Updates or Schema Changes

6.5 View Selectivity

In this experiment, we measure how the performance advantage of incremental view maintenance over recomputation is affected by the view selectivity, i.e., by the probability that a base tuple's data will actually be reflected in a view.

To assess the effect of different view selectivities on view maintenance times, we ran an experiment over different selectivities in the view query. We adjusted selectivity in the range $[0.02 \ldots 1]$ by using different constant values for local conditions in the WHERE-clause of our query (i.e., conditions of the type `FLIGHT.PRICETYPE<=1100`). We define view selectivity over our multiple-relation output schema in analogy to view size as the ratio of the view size of the current query and the view size of a query without WHERE-clause. For each selectivity setting, we deleted a relation with 100 tuples (10% of the input tuples) from the base schema and measured incremental view maintenance time and view recomputation time. Fig. 19 shows the result of the experiment. The graph shows that both incremental view maintenance time and recomputation time increase with the view selectivity, which is not surprising, since in both cases more tuples have to be processed when the view selectivity (and thus the view) becomes larger. However, the relative increase in the incremental update propagation time is similar to the relative increase in recomputation time, meaning that our propagation strategy will keep its performance benefits under changes of the view's selectivity.

7 Related Work

The integration of data stored in heterogeneous schemas has long been an object of intensive studies. The problem of schematic heterogeneity or different source capabilities is repeatedly encountered when attempting to integrate data. Some examples are Garlic [TAH+96], TSIMMIS [HGMI+95], and DISCO [TRV95]. Several logic-based languages have been developed to integrate heterogeneous data sources (e.g., HiLog [CKW89] or SchemaLog [GLS+97]). Some SQL-extensions have also been proposed, such as MSQL [LAZ+89] which has capabilities for basic querying of schema elements, and, in particular, *SchemaSQL* ([LSS96],see below).

Those approaches are able to overcome different classes of schematic heterogeneities. However, the important class of schematic heterogeneities in semantically equivalent relational databases is often excluded from integration language proposals. Miller et al. [MIR93, Mil98] show that relational databases may contain equivalent information in different schemas and give a formal model (Schema Intension Graphs) to determine such "semantic equivalence" of heterogeneous schemas.

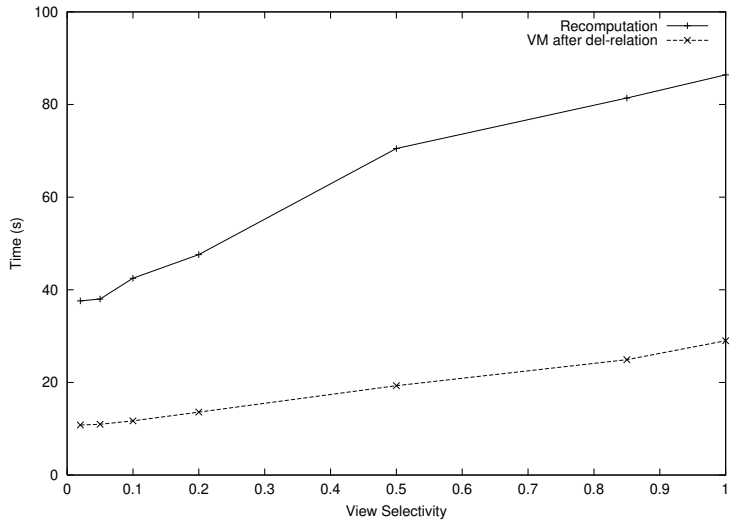

Figure 19: Update Propagation under Views of Different Selectivities

The predominant approach at integrating such semantically equivalent schemas has been done by Gyssens et al. [GLS96] and later by Lakshmanan, Sadri, and Subramanian [LSS96, LSS99]. In [LSS96], the authors present *SchemaSQL*, which is used as the basis for our work. *SchemaSQL* builds upon earlier work in SchemaLog [GLS$^+$97]. It is a direct extension of SQL, with the added capability of querying and restructuring not only data, but also schema in relational databases, and transforming data into schema and vice-versa. Thus, using *SchemaSQL* as a query language makes it possible to overcome syntactic (schematic) heterogeneities between relational data sources.

A second foundation of our work is the large body of work on incremental view maintenance. After early models for view maintenance [BLT86], many algorithms for efficient and correct view maintenance under a variety of assumptions have been proposed. Prominent results, often taking concurrency into account, include ECA [ZGMHW95], SWEEP [AESY97], Mohania et al. [MKK97] and the approach by Gupta et al. [GMS93]. Those approaches follow an *algorithmic* approach in that they propose algorithms to compute changes to a view.

Griffin and Libkin [GL95] consider views with duplicates, and, more importantly, follow an *algebraic* approach which defines a complete and minimal set of relational algebra operators. They achieve a rigorous proof of the correctness of view maintenance by proving the correctness of those operators and their nesting. This is an idea that we have adopted for our work. Griffin and Libkin's work is partly based on the algebraic approach by Qian and Wiederhold [QW91]. Colby et al. [CGL$^+$96] correct the *state bug* that occured in earlier work by Griffin and Libkin as well as in other authors' proposals. Our work builds on existing incremental view maintenance literature but handles schema changes in addition to data updates. Since *SchemaSQL* is able to transform data updates into schema updates and vice-versa, any operator in the *SchemaSQL* algebra tree must be able to propagate both classes of updates. To our knowledge, our view maintenance strategy is the first that achieves incremental update propagation in such schema-restructuring views.

Another body of literature that is related to our work are performance studies on incremental view maintenance algorithms. The work in this field is not as extensive as on view maintenance itself, but a

number of studies exist. An early paper on measuring the performance of incremental view maintenance strategies is Hanson [Han87]. The ECA paper [ZGMHW95] contains a study on the performance of its algorithm, but only in an analytical manner rather than actual performance studies. More recently, there are studies on some view maintenance algorithms for example by O'Gorman et al. [OAE00]. [MKK97] also contains a very simple analysis of their algorithm's efficiency and Griffin and Libkin [GL95] give a analytical complexity study of their algorithm but neither paper evaluates system performance.

8 Conclusions

In this paper, we have proposed the first incremental view maintenance algorithm for schema-restructuring views in *SchemaSQL*. We have shown that the traditional approach at incremental view maintenance—rewriting view queries and executing them against the source data—is not feasible for such views. We have defined an update propagation scheme in which updates are propagated from the leaves to the root in the algebra tree corresponding to the query and proved its correctness. Performance experiments on a prototype of a view-maintenance-capable query engine have shown that update propagation has the expected large benefits over recomputation of views. Lastly, we have also proposed a possible performance improvement by the introduction of *batched-update* primitives. In summary, our work reports a significant step towards supporting the integration of large yet schematically heterogeneous data sources into in integrated environment such as a data warehouse, while allowing for incremental propagation of updates.

Future work includes an extension of our implementation towards such batched updates, a review of the implementation in general with the goal of removing some of the performance obstacles (such as the use of local relations to store intermediate data) and an extension of the implementation towards multi-databases. Another direction of research could be the analysis of the concept presented in this paper with regard to its usability for other data integration languages, such as XML query languages or object-relational languages.

References

[AESY97] D. Agrawal, A. El Abbadi, A. Singh, and T. Yurek. Efficient View Maintenance at Data Warehouses. In *Proceedings of SIGMOD*, pages 417–427, 1997.

[BLT86] J. A. Blakeley, P.-E. Larson, and F. W. Tompa. Efficiently Updating Materialized Views. *Proceedings of SIGMOD*, pages 61–71, 1986.

[CGL+96] L. S. Colby, T. Griffin, L. Libkin, I. S. Mumick, and H. Trickey. Algorithms for Deferred View Maintenance. In *Proceedings of SIGMOD*, pages 469–480, 1996.

[CKW89] W. Chen, M. Kifer, and D. Warren. HiLog as a Platform for Database Languages. *IEEE Data Eng. Bull.*, 12(3):37, September 1989.

[FRV95] D. Florescu, L. Raschid, and P. Valduriez. Using Heterogenous Equivalence for Query Rewriting in Multidatabase Systems. In *3rd Int. Conf. on Cooperative Information Systems*, pages 158–169, Vienna, Austria, 1995.

[GL95] T. Griffin and L. Libkin. Incremental Maintenance of Views with Duplicates. In *Proceedings of SIGMOD*, pages 328–339, 1995.

[GLS96] M. Gyssens, L. V. S. Lakshmanan, and I. N. Subramanian. Tables as a Paradigm for Querying and Restructuring (extended abstract). In *Proceedings of ACM Symposium on Principles of Database Systems*, pages 93–103, 1996.

[GLS+97] F. Gingras, L. V. S. Lakshmanan, I. N. Subramanian, D. Papoulis, and N. Shiri. Languages for Multi-Database Interoperability. In J. M. Peckman, editor, *Proceedings of SIGMOD*, pages 536–538, 1997.

[GMS93] A. Gupta, I. S. Mumick, and V. S. Subrahmanian. Maintaining Views Incrementally. In *Proceedings of SIGMOD*, pages 157–166, 1993.

[GRVB98] J.-R. Gruser, L. Raschid, M. E. Vidal, and L. Bright. Wrapper Generation for Web-Accessible Data Sources. In *6th Int. Conf. on Cooperative Information Systems*, pages 14–23, New York, 1998.

[Han87] E. N. Hanson. A Performance Analysis of View Materialization Strategies. In *Proceedings of SIGMOD*, pages 440–453, 1987.

[HGMI+95] J. Hammer, H. García-Molina, K. Ireland, Y. Papakonstantinou, J. Ullman, and J. Widom. Information Translation, Mediation, and Mosaic-based Browsing in the TSIMMIS System. In *Proceedings of SIGMOD*, pages 483–483, 1995.

[LAZ+89] W. Litwin, A. Abdellatif, A. Zeroual, B. Nicolas, and P. Vigier. MSQL: A Multidatabase Language. *Information Sciences*, 49(1), 1989.

[LSS96] L. V. S. Lakshmanan, F. Sadri, and I. N. Subramanian. SchemaSQL — A Language for Interoperability in Relational Multi-database Systems. In *International Conference on Very Large Data Bases*, pages 239–250, Mumbai, India, Sept. 1996.

[LSS99] L. V. S. Lakshmanan, F. Sadri, and S. N. Subramanian. On Efficiently Implementing SchemaSQL on an SQL Database System. In *International Conference on Very Large Data Bases*, pages 471–482, 1999.

[Mil98] R. J. Miller. Using Schematically Heterogeneous Structures. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 27(2):189–200, 1998.

[MIR93] R. J. Miller, Y. Ioannidis, and R. Ramakrishnan. The Use of Information Capacity in Schema Integration and Translation. In *Proceedings of the Nineteenth International Conference on Very Large Data Bases (VLDB)*, pages 120–133, Dublin, Ireland, August 1993.

[MKK97] M. K. Mohania, S. Konomi, and Y. Kambayashi. Incremental Maintenance of Materialized Views. In *Database and Expert Systems Applications (DEXA)*, pages 551–560, 1997.

[OAE00] K. O'Gorman, D. Agrawal, and A. El Abbadi. On the Importance of Tuning in Incremental View Maintenance: An Experience Case Study (Extended Abstract). In *Proceedings of the Second International Conference on Data Warehousing and Knowledge Discovery (DaWaK)*, pages 77–82, September 2000.

[QW91] X. Qian and G. Wiederhold. Incremental Recomputation of Active Relational Expressions. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 3(3):337–341, September 1991.

[TAH+96] M. Tork Roth, M. Arya, L. M. Haas, M. J. Carey, W. Cody, R. Fagin, P. M. Schwarz, J. Thomas, and E. L. Wimmers. The Garlic Project. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 25(2):557 ff., 1996.

[TRV95] A. Tomasic, L. Raschid, and P. Valduriez. Scaling Heterogeneous Databases and the Design of DISCO. Technical report, INRIA, 1995.

[Ull89] J. Ullman. *Principle of Database and Knowledge-Base Systems.* Computer Science Press, 1989.

[ZGMHW95] Y. Zhuge, H. García-Molina, J. Hammer, and J. Widom. View Maintenance in a Warehousing Environment. In *Proceedings of SIGMOD*, pages 316–327, May 1995.