

# Tiling the view update problem

Zinovy Diskin

Dec 7, 2009

## Abstract

The paper describes a novel algebraic framework for specifying backward translations of view updates. It is based on diagram operations and can be considered an essential generalization of *dynamic* views (known in the literature since the eighties). A distinctive feature of the framework is that both view and update mappings are first class citizens and explicitly occur in the arity shape of the translation operation. We show that dynamic views augmented with mappings give rise to double categories, and discuss the benefits of this formalism.

## 1 Introduction

The *view update problem* refers to the task of translating updates to a view to updates on the database. Such translations are *correct* if the view of the updated base equals to the updated view. Since the view abstracts away some information, many (or none) correct updates on the base may result in the given view update. Finding a unique backward translation is a difficult problem, which does not have universal solutions. Yet for special classes of queries and updates, reasonable translation procedures ensuring uniqueness (so called *update policies*) could be found. Design of update policies is a delicate business that needs a careful inspection of the syntactical (the base schema and the view definition) and the semantic (conditions on instances) sides of the problem [4, 12, 14].

To manage complexity, it is reasonable to split the task into simpler components. A complex view definition  $V$  is decomposed into a sequence of simple blocks  $\Delta_i V$  (so that  $V$  becomes a view of a view of...), and similarly a complex update  $U$  is presented as a sequence of elementary deletions, insertions and modifications  $\Delta_j U$ . For every pair  $(\Delta_i V, \Delta_j U)$  a correct translator  $T_{ij}$  is a priori designed and stored in a library. Then translation  $T(U)$  of the update  $U$  over the view definition  $V$  can be composed from elementary blocks  $T_{ij}$  as schematically shown in Fig. 1(a). This mechanism works if we can guarantee correctness of  $T(U)$  as soon as all “tiles”  $T_{ij}$  are correct.




Figure 1: Tiling update translation

In overall, the goal of the paper is to make the sketch above mathematically precise: to explicate the underlying algebraic structures and show their adequacy to the problem.

In more detail, the tiling idea separates the view update problem into two parts. The hard one is to create the very building blocks (tiles), i.e., to design update policies and write translators  $T_{ij}$  — we are not dealing with it in this paper. The other part is to design a composition mechanism for joining elementary translators together (tiling). We need to specify a common interface for tiles  $T_{ij}$  so that tiling would be implementation-independent and correctness-preserving. Although easier, this part is not straightforward and exposes several issues to be solved.

As mentioned above, update policies are inherently specific and dependent on particular kinds of schema constraints (dependencies), queries, updates, and their interaction. The tile interface should capture peculiarities of individual policies in an abstract way. On the other hand, the interface should *not* be overly abstract and skip essential constructs because it may corrupt tile composition (tiling must be “tight”). Note also that the very composition is two-dimensional, which requires special algebraic means beyond ordinary composition of functions. We also aim to formulate tiling mechanism in a syntax-independent way to make it applicable for models beyond relational. Indeed, “viewOf” relationships are ubiquitous in software engineering and different ver-

sions of the view update problem appear in different contexts beyond the traditional database area, for example, in bi-directional programming [7, 18], optimistic replication [8] and model synchronization [17, 5].

Use of algebra and compositionality is, of course, a known idea that can be traced back to the first works on the view update problem. Compositionality of updates is a basic assumption in founding works on constant complement [1] and dynamic views [10]. Compositionality of views is an evident requirement, but it has not been explicitly formulated until the recent work on lenses [2]: sequential lens composition [7] includes a trivial discrete treatment of view composition. Unfortunately, algebraic structures used in these works are too rough for our goals: they model the view and base datasets as discrete points with nothing in-between. That is, in terms of Fig. 1(a), arrows between nodes have no extension: they are just pairs of nodes and denote “viewOf” and “updateOf” relationships. The reality is richer. If a dataset  $B$  is a view of  $A$ , there is a *view mapping*  $v: B \rightarrow A$  tracing the origin of  $B$ ’s elements. Similarly, if a dataset  $A'$  is an update of  $A$ , there is an *update mapping*  $u: A \rightarrow A'$  specifying which elements were kept unchanged and which ones were modified in the update. In section 2 we consider an example justifying the claim. Thus, arrows in Fig. 1(a) do have extension and the tile composition mechanism must ensure their coherence. Figure 1(b) shows that translations well composable in the discrete framework may be not composable if mappings are taken into account.

Another key point to be addressed is syntax. To broaden applicability and avoid peculiarities of a particular syntactic mechanism, both lenses and dynamic views only consider views semantically, i.e., as functions from one space of datasets to another. We will say these frameworks are *syntax-free*. However, policy design does account for typing: data elements are typed by schema elements, view traceability links are typed by view definitions and update links are typed by update definitions ([4] is a notable example of how it works). Thus, in an accurate specification of update translation, syntax must be captured but in an abstract way. The framework is to be *syntax-independent* rather than syntax-free.

To summarize, we aim to specify a mathematical structure satisfying the following requirements:

(a) Be concrete enough to capture the following constructs essential for the view update problem: data, schemas and typing, view and update definitions, view and update mappings, and update translation

all working in concert;

(b) Be abstract enough to model these constructs in a syntax-independent way;

(c) Address compositions of views and of updates, and their interaction between themselves and with update translations;

(e) Be algebraic to continue the “product line” of algebraic models for the view update problem.

Appearance of mappings in the requirement (a) and especially requirement (c) make use of category theory practically inevitable. Our main algebraic vehicle is the notion of double category, which is nothing but a systematic algebraic treatment of two-dimensional composition. In the course of the paper we explain categorical ideas in parallel with their application.

Our plan is as follows. In section 2 we consider a detailed example of relational view update to show that mappings do matter. We also demonstrate the deficiency of fixing one predefined update policy, and argue that a possibility to choose from a range of policies works better (sections 2.3,2.4). In section 3 we specify a mathematical structure satisfying the requirements; we try to reduce categorical “abstract nonsense” to a minimum and focus on ideas rather than technicalities.

## 2 Mappings and view updates

Figure 2 presents an example of relational update propagation, inspired by the main example of [4].<sup>1</sup> Our goal in this section is first to illustrate the main concepts and constructs of the view update translation, and then abstract them in semi-formal terms to prepare and motivate the subsequent syntax-independent formalization (in section 3). Each subsection thus begins with considering a corresponding fragment of the example, and then proceeds to its semi-formal abstract arrangement.

Distinction between a mapping (function) and its instance (a link or tuple) is crucial for this section. Given a mapping  $f: X \Rightarrow Y$  between sets or structures  $X, Y$ , we write the value of  $f$  at argument  $x \in X$  as either  $f(x)$  or  $(x)f$  or  $x.f$  choosing notation that makes formulas easier to read. If  $y = f(x)$ , the pair  $(x, y)$  is called an (*application*) *instance* of  $f$ . In instance diagrams like Fig. 2, we show such pairs by arrows  $x \rightarrow y$  and call them *links*. Thus, a mapping  $f$  appears as a set of links, and in instance diagrams we show mappings by ovals encompassing their links. In

<sup>1</sup>We refer to [4] for science-historical explanations and possible complaints :).

*ref to lechten-  
borger pods pa-  
per is lost*

the text, we distinguish mappings and links by using double- and single-line arrows respectively.

## 2.1 View definition

Consider three tables in the left top corner of Fig. 2. The database *schema*  $\mathcal{S}$  consists of two relation *schemes*, EDS(Empl,Dept,Sal) and DM(Dept,Mngr), and its possible state  $A$  is presented by the corresponding tables (with an intended typo in the Newton’s name). The columns of table EDS are displayed in the “reverse order” to reduce clutter later when we will show mappings between relations. The names of key columns are underlined. The bullets in the columns under relation names denote row identifiers (rids), which will be referred to as EDS.row1, EDS.row2 etc. counting from the top. The set of rids of table  $T$  is denoted by  $\llbracket T \rrbracket^{rid}$  and called the *rid extension* of  $T$ . The extension of the entire table (all columns including rid) is denoted by  $\llbracket T \rrbracket$ .

Table EDS $\otimes$ DM is derived by the following query  $Q$ :

$$EDS\otimes DM \stackrel{\text{def}}{=} (EDS \otimes_{Dept=Dept} DM) \upharpoonright_{(Empl,Mngr)}$$

where  $\otimes_{-}$  denotes the relational join operator (subindexed by the condition) and  $\upharpoonright_{-}$  denotes projection to the set of attributes in the subindex. More accurately, the syntactical side of the table (the attribute and relation names, metadata) is not derived and is a part of the query definition. The very data in the table are computed/derived by query execution, correspondingly, the respective part of the table is framed with dashed lines (blue with a color display). Rids are also derived: each one is a pair of rids from the respective operands, (EDS $\otimes$ DM).row1=(EDS.row1,DM.row1) and (EDS $\otimes$ DM).row2=(EDS.row2,DM.row2), which is also shown by dashed blue projection arrows. In general, the rid extension of a join table is a subset of the Cartesian product of the operands,  $\llbracket T_1 \otimes_P T_2 \rrbracket^{rid} \subset \llbracket T_1 \rrbracket^{rid} \times \llbracket T_2 \rrbracket^{rid}$  with the subset selection condition given by predicate  $P$ . We thus have two projection mappings  $p_i: \llbracket T_1 \otimes T_2 \rrbracket^{rid} \Rightarrow \llbracket T_i \rrbracket^{rid}$ ,  $i = 1, 2$ .

A view to  $\mathcal{S}$  is specified in the right top corner of the figure. The view schema  $\mathcal{T}$  consists of a single relational scheme SM(Scholar,Master), which is mapped to the base schema  $\mathcal{S}$  by three bent links at the very top of the figure. These three links form a mapping  $v: \mathcal{T} \Rightarrow Q(\mathcal{S})$ , where  $Q(\mathcal{S})$  denotes the union of  $\mathcal{S}$  with the query schema  $\mathcal{S}_Q$  (in our example,  $Q(\mathcal{S})$  consists of three relation schemes EDS, DM, EDS $\otimes$ DM). We need the union rather than only

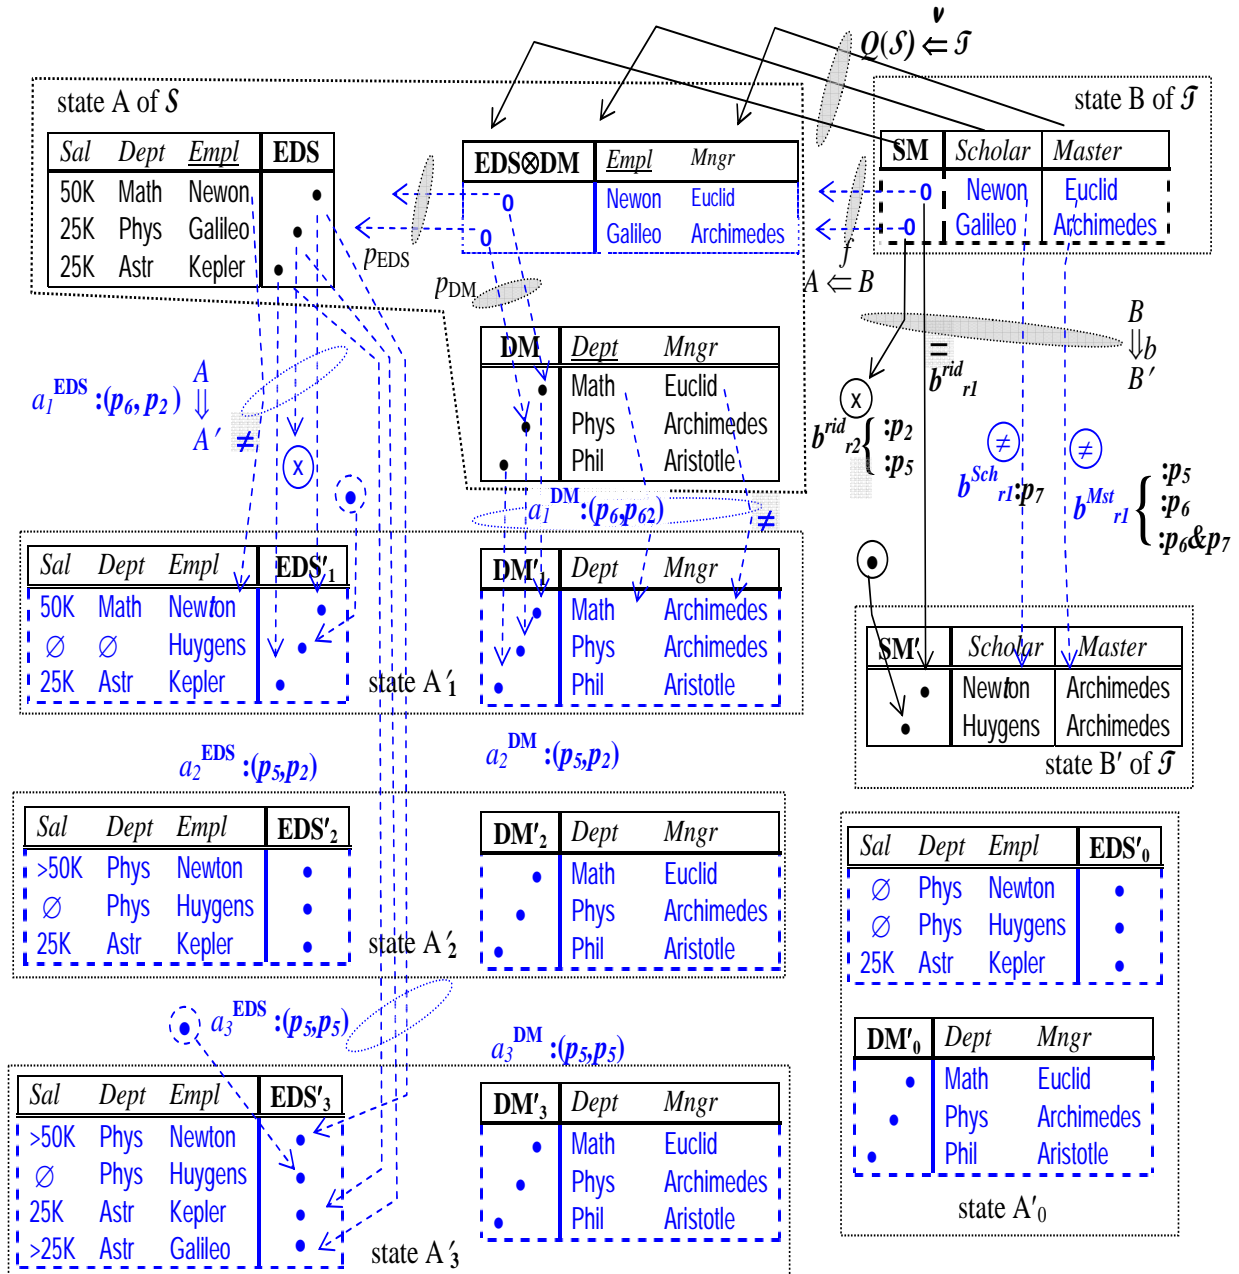
$\mathcal{S}_Q$  because in general some of the view schema relations are mapped to basic relations of  $\mathcal{S}$ . Thus, a view definition is a triple  $(\mathcal{T}, Q, v)$  with  $\mathcal{T}$  the view schema,  $Q$  a query against  $\mathcal{S}$ , and  $v$  a mapping as above. Importantly, the view mapping must respect the constraints, e.g., key columns are mapped to key columns.

Given a view definition, the extension of the view schema is computed by, first, executing the query (and getting table EDS $\otimes$ DM) and then coping the result to the view location (the table SM). More accurately, first an isomorphic set of fresh rids is created and then data in the attribute columns are copied. We thus have a mapping  $f: \llbracket SM \rrbracket \Rightarrow \llbracket EDS\otimes DM \rrbracket$ , whose rid part is shown in the figure.

## 2.2 Updates and update propagation

Let  $B$  denotes the initial state of the view shown in the figure. Suppose that the view is updated and comes to a state  $B'$  (table SM'). Let us first ignore all arrows in-between the two states and inspect the states immediately. We see that two old rows were deleted and two new rows are inserted. Deletions can be propagated to the database by deleting the respective rows in either one or both of the base relations EDS, DM (or by their modifications destroying the joining condition). Assume, for example, that deletions in a join table are realized by suitable deletions in the left operand of the term  $T_1 \otimes T_2$ , (cf.[2]), i.e., table EDS in our case. We also assume that insertions into SM are realized by insertions into EDS, and so we add two new rows to EDS with Nulls in the Salary column denoting unknown values. The resulting database state  $A'_0$  is shown on the right of the figure under the table SM'.

However, here is another interpretation of the passage from SM to SM'. Suppose that the row1 was modified rather than deleted: the typo in the name Newon was fixed, and Newton’s Master changed from Euclid to Archimedes. We show this interpretation by linking the rids: arrow  $b_{r_1}^{rid}$  linking SM.row1 and SM'.row1 means that two rows refer to the same real world object (note the =-label near the link). The absence of such linking for SM.row2 is syntactically sugared by a quasi-link  $b_{r_2}^{rid}$  from the row to the “crash-symbol”  $\otimes$ . Similarly, the absence of links going to SM'.row2 is sugared by a “creation” link from symbol  $\odot$ . Due to rigidity of the relational model, links for the rids determine the corresponding links for the values in the attribute columns, which may be immediately tested for equality and we get  $\neq$ -links from Newon to Newton and from Euclid to Archimedes.



Update Translation Table							
Action over $\mathcal{T}$ (i.e., $SM$ )	Action over $\mathcal{S}$						
	Hire Empl.	Fire Empl.	Create Dept.	Delete Dept.	ChangeOf Dept.	ChangeOf Mngr.	Modify Attr.
	$p_1$	$p_2$	$p_3$	$p_4$	$p_5$	$p_6$	$p_7$
Delete row ( $u_1$ )	-	+	-	+	+(moveToNomng.)	+(Del Mngr.)	-
Insert row ( $u_2$ )	+	-	+	-	+(moveFromNomng.)	+(Get Mngr.)	-
Modify <i>Scholar</i> clmn. ( $u_3$ )	-	-	-	-	-	-	+
Modify <i>Master</i> clmn. ( $u_4$ )	-	-	-	-	+	+	-

Figure 2: Examples of relational view updates: the base schema  $\mathcal{S} = \{EDS, DM\}$  and the view schema  $\mathcal{T} = \{SM\}$

In this way we obtain an *update mapping*  $b: B \Rightarrow B'$  consisting of links  $b_{ri}^x$ ,  $i = 1, 2$ ,  $x = rid, Sch, Mst$ . We will see in the next subsection that this interpretation of the update has several translation strategies that result in states different from  $A'_0$ .

We can imagine yet another update mapping consisting of the link  $b_{r1}^{rid}$  as above, and the link from SM.row2 to SM'.row2: think of the case when Galileo was erroneously placed into Physics instead of Huygens. Backward translation of this update would result in table EDS' different from EDS'\_0 by the value 25K rather than  $\emptyset$  in Salary column. Thus, update mappings do matter.

Note that in general update mappings cannot be derived from the states. If we attempt to identify rows automatically by primary keys, then it would be impossible to modify the key (e.g., to fix a typo). If, on the contrary, we attempt to identify rows automatically by internal identifiers (rids), it would not work too because the user can delete a row but later recognize that was a mistake and restore the same row — however, the system would assign this same row a new identifier. Automatic rid-identification would not also work when the view is copied to another computer. Thus, update mappings is an independent part of update specification.

Update mappings reduce uncertainty in interpreting updates but not entirely. For example, there is a number of different database changes that cause the change of Newton's Master from Euclid to Archimedes in the view, and a number of different changes that cause deletion of the row (Galileo, Archimedes). We first fix one possible interpretation (cause 1) to see how the translation works, and later will address variability (causes 2,3,...).

Suppose that update  $b: B \Rightarrow B'$  shown in Fig. 2 is caused by the change of the Math manager from Euclid to Archimedes and by firing Galileo. The corresponding base update is specified by mapping  $a_1: A \Rightarrow A'_1$ , whose components,  $a_1^{EDS}: A^{EDS} \Rightarrow A'^{EDS}_1$  and  $a_1^{DM}: A^{DM} \Rightarrow A'^{DM}_1$ , are shown in the figure (some links are missing to avoid clutter but we remind that a complete set of rid-links determines the rest. Note that Huygens' department is unknown because now Archimedes manages both Math and Physics). These two component are actually a projection of a bigger mapping  $Q(a_1): Q(A) \Rightarrow Q(A'_1)$  including also the third component  $a_1^{EMS \otimes DM}: EMS \otimes DM \Rightarrow EMS' \otimes DM'$  not shown in the figure.

The entire procedure of view computation and backward update propagation is specified in Fig. 3

in abstract terms of objects (data structures) and mappings. Horizontal and vertical arrows denote mappings related to views and updates respectively. “Depth” arrows going from the front to the back side are typing mappings: they provide types (relation and attribute names) for data elements. The user-defined elements are shown by black solid lines, the computed elements are dashed (blue with a color display); basic nodes are shaded and derived ones are blank. Computation consists of five consecutive operations (denoted by double-headed arrows with numbers), which input two arrows of a certain incidence, and output other two arrows that extend the input configuration up to a square. Substituting results of one operation to another amounts to tiling. This is a key advantage of the notation: substitution becomes geometrically evident.

The operations (steps) are as follows.

*Step 1<sup>1</sup>. Query execution.* A query against schema  $\mathcal{S}$  is specified by adding to  $\mathcal{S}$  new derived elements to be computed by the query. Hence, we have an inclusion mapping  $i_Q: \mathcal{S} \hookrightarrow Q(\mathcal{S})$ . A database state is a typing mapping  $\tau_A: A \Rightarrow \mathcal{S}$  (see the top horizontal left square in the figure). Given a query definition  $i_Q$  and data  $\tau_A$ , the query can be executed; adding the execution results to the initial dataset can be specified by inclusion  $i_Q^A: A \hookrightarrow Q(A)$ .

*Step 1<sup>2</sup>. Coping and relabeling.* The results of the query are copied to the view according to the view mapping  $\mathbf{v}$ .

*Step 2<sup>1</sup>. Update execution.* An update request  $\mathbf{u}$  is submitted to the view and executed producing an update mapping  $b$  (the rightmost face of the parallelepiped).

*Step 2<sup>2</sup>. Backward relabeling.* The update results are relabeled according to the augmented schema  $Q(\mathcal{S})$  (and may be thought of as “copied” back to virtual slots storing the query results). We may also consider this operation as backward projection of  $B'$  along  $f$ , hence the projection notation  $\_ \downarrow_f$  for the results.

*Step 3. Backward propagation.* The query schema extension is backward propagated to the database. This propagation can be considered as finding solution  $A'$  to algebraic equation  $Q(A') = B' \downarrow_f$ . As in general there are many solutions, to achieve uniqueness some update policy is needed. This is the most non-trivial part of the procedure.

Diagram Fig. 2(b) presents an abbreviated version of diagram (a): we hide the query mechanism inside the view mappings  $\mathbf{v}_Q$ ,  $f_Q$  and operation  $\mathbf{vExe}$  of view execution which includes query execution. The idea is made precise in categorical algebra, and is known



$x=rid, Sch, Mst$  constituting the update  $b: B \Rightarrow B'$  is labeled with a set of possible database update  $\mathbf{p}_i$  causing this link to appear. By choosing one of the labels from the set, we set a unique backward interpretation of the view update, for example, taking  $\mathbf{p}_5$  for  $b_{r1}^{Mst}$  and  $\mathbf{p}_2$  for  $b_{r2}^{rid}$  states that Newton moved to the department managed by Archimedes and Galileo quitted. It does not provide a unique database update (if there are several departments managed by Archimedes), but essentially narrows the range of possibilities. This range can be described by conditional values, for example, we place the value “a dept. managed by Archimedes” in the column  $EDS_2Dept$  for row Newton. In our case, this value is uniquely evaluated to Physics and we come to update  $a_2: A \Rightarrow A_2$  (note its label  $(\mathbf{p}_5, \mathbf{p}_2)$ ). Thus, we have six possibilities of backward translations (six combinations of labels for the links), three of which are shown in the figure. In general, the same update  $b: B \Rightarrow B'$  may have a number of quite reasonable and correct backward interpretations  $a: A \Rightarrow A'$ .

## 2.4 ... and changing the focus.

The conclusion above is well-known, and it is what constitutes the view update problem. Our goal in the paper is to provide general and unifying mathematical foundations for different approaches to the problem. To this end, we consider the most general situation when the view user is allowed to label view updates with possible database updates to inform the system about the user’s interpretation (or intention) with the update. To subsume the case of a predefined update policy, we may introduce a special null label “no interpretation”, in which case the system translates the update according to some predefined “null” policy.<sup>2</sup>

Thus, a view update request becomes a pair  $(\mathbf{u}_i, \mathbf{p}_j)$  with  $\mathbf{u}_i$  the view update definition and  $\mathbf{p}_j$  its intended backward interpretation. This settings changes the focus of the view update problem. The main task is the design of update definition spaces,  $U_{\mathcal{S}}$  and  $U_{\mathcal{T}}$ , and enumeration of reasonable correspondences between them (the Translation Table of our example). The previous focus becomes a particular branch related to the “null”-label, that is, to backward translation of pairs  $(\mathbf{u}, \emptyset) \in U_{\mathcal{S}} \times U_{\mathcal{T}}$ . The approach outlined above seems appealing but its precise mathematical underpinning is not straightforward. Con-

<sup>2</sup>By endowing the system with inductive learning and other AI capabilities, we may gradually make the backward translation more reasonable by using heuristics, or/and modifying the interface for update request.

sidering view updates as pairs  $(\mathbf{u}_i, \mathbf{p}_j)$  means that the translation mapping is a binary relation  $\mathbf{T} \subset U_{\mathcal{T}} \times U_{\mathcal{S}}$  rather than a dynamic view function  $\mathbf{T}: U_{\mathcal{T}} \Rightarrow U_{\mathcal{S}}$  [10]. Moreover, our example shows that the update spaces are graphs of states and update mappings rather than plain sets of states. Thus,  $\mathbf{T}$  is a binary relation between graphs. In addition, we have seen that the spaces are interconnected by view mappings  $f: B \Rightarrow A$ . Finally, both update and view mappings are composable and translation should be compatible with these compositions in some way. In the next section we specify a suitable mathematical framework.

## 3 Well-defined view systems

We begin with a short primer on graphs and categories. To ease reading, we will often write “a set  $X$  of *widgets*” instead of “an abstract set  $X$  whose elements are called *widgets*”.

### 3.1 From graphs to double categories

**Graphs.** A *graph*  $G$  consists of a set of *nodes*  $G_0$  and a set of *arrows*  $G_1$  together with two mappings  $\partial_i: G_1 \rightarrow G_0, i = 0, 1$ . As usual, we write  $a: N \rightarrow N'$  if  $\partial_0 a = N$  and  $\partial_1 a = N'$ . Sometimes we write  $a: N \rightarrow N'::G$  to remind that  $a \in G_1$ . Two arrows  $a, b$  are called *composable* if  $a\partial_1 = \partial_0 b$ . A graph is called *thin* if for any pair of nodes  $(N, N')$  there is at most one arrow  $a: N \rightarrow N'$ . A *graph mapping* (*morphism*)  $f: G \rightarrow G'$  is a pair  $(f_0, f_1)$  of mappings (functions)  $f_i: G_i \rightarrow G'_i$  that preserves the incidence relations between nodes and arrows ( $a.f_1.\partial_0 = a.\partial_0.f_0$ , etc.).

**Tiles.** Let  $\square$  denotes a fixed graph shown in Fig. 4(a) with fixed names for its nodes and arrows. A *tile* in a graph  $G$  is a graph mapping  $T: \square \rightarrow G$ . It is convenient to denote it as shown in Fig. 4(b) (ignore symbol  $\mathbb{P}$  for a moment), which means that  $T(12) = f, T(13) = b, T(34) = f',$  and  $T(24) = a$ . In other words, a tile is a quadruple of arrows  $(f, b, f', a)$  subject to incidence relationships shown in diagram (b) and encoded by diagram (a). If we write a tile as a quadruple of arrows, we always follow the order used above (12,13,34,24). Some nodes and some arrows in a tile may be equal (it’s a tuple, not a set), e.g., it may happen that  $T(2)=T(4) = A$ , in which case both left nodes in diagram (b) will be  $A$ . Nodes  $A, B..$  and arrows  $f, b..$  are *elements* of the tile  $T$ . They will be denoted by  $\bullet T, T\bullet, ..$  and  $\bar{T}, T|, ..$  respectively. If  $\mathbb{P}$  is a class of tiles (a predicate that defines the class), we write  $T:\mathbb{P}$  for  $T \in \mathbb{P}$  and place this declaration at the center of the diagram. The name of the tile may




Figure 4: Tile (a,b) and db-arrow (c) definitions

be left implicit but the predicate is always there.

**Categories and functors.** A *category*  $\mathcal{C}$  is a graph with

- (i) an associative operation of arrow composition denoted by  $;$  (semi-colon), that is, any pair of composable arrows  $A \xrightarrow{a} B \xrightarrow{b} C$  is assigned with one and only one arrow  $A \xrightarrow{c} C$  denoted by  $a;b$ , and
- (ii) a unary operation  $\mathbf{1}_$  that assigns to every node  $N$  an *identity/idle* loop  $\mathbf{1}_N: N \rightarrow N$ . These loops are units of the composition:  $\mathbf{1}_{\partial_0 a}; a = a = a; \mathbf{1}_{\partial_1 a}$ . Nodes in a category are called *objects*, and arrows are *morphisms*. The graph underlying category  $\mathcal{C}$  is denoted by  $|\mathcal{C}|$ , but we will often omit the bars. Hence,  $\mathcal{C}_0$  and  $\mathcal{C}_1$  denote the classes of all objects and all morphisms resp. The former can be also considered as a category, whose only arrows are identities (called *discrete category*).

A *functor*  $f: \mathcal{C} \rightarrow \mathcal{C}'$  is a graph morphism  $f: |\mathcal{C}| \rightarrow |\mathcal{C}'|$  that is compatible with arrow composition  $(a;b).f_1 = (a.f_1); (b.f_1)$  and idle loops  $f_1(\mathbf{1}_N) = \mathbf{1}_{N.f_0}$ .

**Double graphs, categories, functors.** A *double graph*  $\mathbb{D}$  is a four-sorted algebraic structure comprising a class  $\mathbb{D}_0$  of *nodes* (also called *0-cells*), classes  $\mathbb{D}_1^h$  of *horizontal* and  $\mathbb{D}_1^v$  of *vertical arrows* (*1-cells*) and a class  $\mathbb{D}_2$  of *double arrows* (*2-cells*) subject to incidence (source and target) conditions shown in diagram Fig. 4(c). To wit: if  $\Delta$  denotes a double arrow, then its *vertical source* is arrow  $f$ , the *horizontal target* is  $a$  etc. These data define a tile  $\partial\Delta$  called the *boundary* of  $\Delta$ . Below we use abbreviations *h-arrow*, *v-arrow*, *db-arrow* or *db-graph*.

A db-graph is called *thin* if for any tile  $T$  there is at most one db-arrow  $\Delta$  such that  $\partial\Delta = T$ .

A *double category* is db-graph satisfying the following requirements. Nodes and h-arrows form a category denoted, again, by  $\mathbb{D}_1^h$ , and similarly, we have a category  $\mathbb{D}_1^v$  so that  $(\mathbb{D}_1^h)_0 = (\mathbb{D}_1^v)_0 = \mathbb{D}_0$ . H-arrows considered as *nodes* and db-arrows as *arrows* form a category  $\mathbb{D}_2^v$ , in which db-arrows are composed *vertically*. similarly, v-arrows as *nodes* and db-arrows as *arrows* form category  $\mathbb{D}_2^h$  in which db-arrows are

composed *horizontally*. Thus, both categories have the same set of arrows,  $(\mathbb{D}_2^h)_1 = (\mathbb{D}_2^v)_1 = \mathbb{D}_2$ , but compositions are different. Nevertheless we denote them by the same symbol, semi-colon, because the context is always clear.

The fact that  $\mathbb{D}_2^x$ ,  $x=h,v$ , are categories means that there are *h-idle* and *v-idle* db-arrows, which are the units of the respective compositions. Horizontal edges of h-idle db-arrows must be idle h-arrows; same holds for v-idle db-arrows. (These definitions become quite clear when the diagrams are drawn on a paper, see [13] or [6]). Finally, horizontal and vertical db-arrow compositions are coordinated between themselves by an *interchange law*: for a regular net of db-arrows like in Fig. 1(a), any order of composition (first horizontally and then vertically, or the other way round, or in a mixed way) gives the same result (see [13] for details).<sup>3</sup>

A *double functor*  $f: \mathbb{D} \rightarrow \mathbb{E}$  is a quadruple of functions  $f_0, f_1^v, f_1^h, f_2$ , sending  $i$ -cells in  $\mathbb{D}$  to respective  $i$ -cells in  $\mathbb{E}$  ( $i = 0, 1^v, 1^h, 2$ ), with preservation of all incidence relationships and such that pairs  $(f_0, f_1^x): \mathbb{D}_1^x \rightarrow \mathbb{E}_1^x$  and  $(f_1^x, f_2): \mathbb{D}_2^x \rightarrow \mathbb{E}_2^x$  ( $x=v,h$ ) are ordinary functors.

## 3.2 Building the definition

We aim to define a mathematical structure satisfying the list of requirements described in Introduction. We approach the task with a sequence of steps that consecutively introduce new “pieces” of structure until the requirements (a,c) are satisfied.

### 3.2.1 The universe

We take some infinite directed graph  $\mathcal{U}$  to be the carrier of the mathematical structure we are going to build. Nodes of  $\mathcal{U}$  are interpreted as objects with structure of some kind (e.g., graphs with additional structure). Arrows of  $\mathcal{U}$  are interpreted as mappings between these objects compatible with their structure. We will use terms *node* and *object*, *arrow* and *mapping* interchangeably.

Some classes of  $\mathcal{U}$ -object may have additional structure, and some classes of mappings may respect this additional structure. We will thus have several sorts of objects and mappings. Some pairs of composable mappings can be indeed composed into a “long” arrow while others cannot depending on sorting. We say that  $\mathcal{U}$  is an infinite graph with *partially composable* arrows.

<sup>3</sup>Appendix reproduces this material for the reviewers’ convenience but is excluded from the paper due to space limitations.



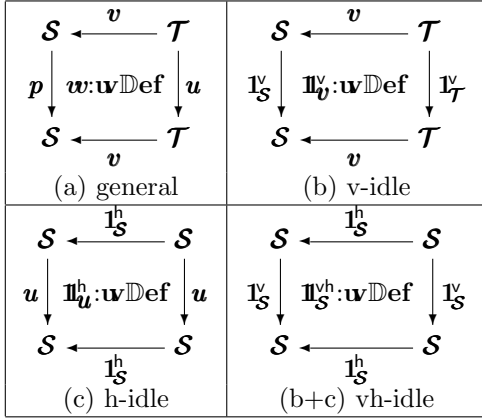


Figure 5: Class of  $\mathbf{wDef}$ -tiles

### 3.2.2 Metadata

**A: Data, view and update definitions.** There is a subclass *dataDef* of  $\mathcal{U}$ -objects called *data definitions* or *schemas*. There are two sorts of arrows between schemas: *update definitions* and *view definitions* (the former are loops). When talking about metadata, we will shortly say updates and views. Each schema  $\mathcal{S}$  is assigned with two special loop-arrows: *idle update*  $\mathbf{1}_{\mathcal{S}}^{\mathcal{S}}: \mathcal{S} \rightarrow \mathcal{S}$  (that defines an idle action doing nothing) and *identical view*  $\mathbf{1}_{\mathcal{S}}^{\mathcal{S}}: \mathcal{S} \rightarrow \mathcal{S}$  (that defines  $\mathcal{S}$  as a view of itself). Updates can be composed associatively, and the idle update is neutral w.r.t. composition. It means that schemas and updates form a category, *updDef*. Similarly, we postulate a category of schemas and views *viewDef*. These two categories have the same class of nodes but different arrows. Together they form a subgraph  $\mathcal{U}^{\text{def}}$  of  $\mathcal{U}$ . We additionally assume that each schema  $\mathcal{S}$  has a designated set of basic updates  $\mathbf{U}_{\mathcal{S}}$  so that all  $\mathcal{S}$ -updates are freely generated by  $\mathbf{U}_{\mathcal{S}}$ . To ease reading the diagrams below, we will always draw view arrows horizontally and updates vertically.

**B: Forward update translation.** An ordinary view definition specifies how to transform data. We consider an extended notion of view, which specifies also how to transform updates.

We assume defined a class  $\mathbf{wDef}$  of tiles of shape specified in Fig. 5(a). It is interpreted as that update  $\mathbf{p}$  on the base is translated to update  $\mathbf{u}$  on the view. Note that two horizontal arrows are equal: translation does not change the view definition.

Class  $\mathbf{wDef}$  satisfies the following conditions.

(o) We assume that any basic view update  $\mathbf{u}$  is correct and hence there is at least one  $\mathbf{wDef}$ -tile having  $\mathbf{u}$  as its right edge.




Figure 6:  $\mathbf{wDef}$ -tile composition

(a) For any update  $\mathbf{p}$ , there is one and only one  $\mathbf{wDef}$ -tile. Hence, we may consider  $\mathbf{wDef}$  as an operation producing arrow  $\mathbf{u}$  from the other elements of the tile. Condition (o) states that this operation is surjective.

(b) The class contains special tiles specified in diagrams Fig. 5(b,c). That is, an idle update is forward translated into the corresponding idle update, and identical views do nothing with updates. It implies that all tiles consisting of four idle arrows belong to  $\mathbf{wDef}$  (diagram (b+c)).

(c)  $\mathbf{wDef}$ -tiles can be composed horizontally and vertically as shown in Fig. 6. In more detail, we define outer tiles  $\mathbf{w}_{32}; \mathbf{w}_{21} \stackrel{\text{def}}{=} [(v_{32}; v_{21}), \mathbf{u}_3, (v_{32}; v_{21}), \mathbf{u}_1]$  and  $\mathbf{w}_{21}; \mathbf{w}'_{21} \stackrel{\text{def}}{=} [v_{21}, (\mathbf{u}_2; \mathbf{u}'_2), v_{21}, (\mathbf{u}_1; \mathbf{u}'_1)]$ , and the condition states they belong to  $\mathbf{wDef}$  as soon as the inner tiles are such.

Below, names of definitions (data, update, view) are bold italic.

### 3.2.3 Data vs. metadata

For any schema  $\mathcal{S}$ , there is a class of arrows into  $\mathcal{S}$  disjoint to  $\mathcal{U}^{\text{def}}$  and called *typing mappings*. A pair  $A = (D_A, \tau_A)$  with  $\tau_A: D_A \rightarrow \mathcal{S}$  a typing mapping is to be thought of as *data object*  $D_A$  schematized by  $\mathcal{S}$ . We also call such pairs *datasets* over  $\mathcal{S}$ , or (*data*) *instances* of  $\mathcal{S}$ , or (*database*) *states* when  $\mathcal{S}$  is considered a database schema. The class of  $\mathcal{S}$ -instances is denoted by  $\llbracket \mathcal{S} \rrbracket$ . We will also write  $A:\mathcal{S}$  for  $A \in \mathcal{S}$ .

### 3.2.4 Execution mechanisms

**A1: Update execution.** We assume defined some class  $\mathbf{uExe}$  of tiles of the shape shown in Fig. 7(a): the left arrow is an update definition, the horizontal arrows are typing mappings, and arrow  $\mathbf{a}$  is to be thought of as execution of  $\mathbf{u}$  for dataset  $A=(D_A, \tau_A)$ .

Class  $\mathbf{uExe}$  satisfies the following conditions.

(a) For every arrow  $\mathbf{p}: \mathcal{S} \rightarrow \mathcal{S}$  and data instance  $A = (D_A, \tau_A)$  of  $\mathcal{S}$ , there one and only one  $\mathbf{uExe}$ -tile




Figure 7: Update execution

having  $\mathbf{p}$  and  $\tau_A$  at its left top corner as shown in diagram (a). Hence, we may consider  $\mathbf{uExe}$  as an operation producing instance  $A'$  and vertical arrow  $a$  from the other elements of the tile. We denote instance  $A'$  by  $A^{\mathbf{p}}$ , arrow  $a$  by  $\overline{\mathbf{p}}_A$ , and the entire tile by  $[\mathbf{p}]_A$ . Arrows like  $a$  are called *update mappings*

(b) All tiles with idle update definitions Fig. 7(b) are in  $\mathbf{uExe}$ . In this way a subclass of *idle* update mappings denoted by  $\mathbf{1}_A$  is defined.

(c) Class  $\mathbf{uExe}$  is closed under vertical tile composition Fig. 7(c). That is, if two inner tiles are from  $\mathbf{uExe}$ , composition  $a; a'$  is defined in  $\mathcal{U}$  and the outer tile is also in  $\mathbf{uExe}$ . We assume that composition so defined is associative, and tiles specified by diagram (b) are units. The latter means that idle update definition produce idle update mappings. Hence, our assumptions can be summarized by saying that instances as objects and  $\mathbf{uExe}$ -tiles as arrows between them form a category. We will denote it by  $\mathbf{updMap}$ .

**A2: View execution.** We assume defined some class  $\mathbf{vExe}$  of tiles of the shape shown in Fig. 8(a). The upper horizontal arrow is a view definition, the vertical arrows are typing mappings—the right one is computed by executing the view definition for dataset  $A=(D_A, \tau_A)$ , the bottom arrow is to be thought of as the view traceability mapping. We denote the entire tile by  $[\mathbf{v}]_A$  and write  $A^{\mathbf{v}}$  for  $B$  and  $\overline{\mathbf{v}}_A$  for *view mapping*  $f$ .

Class  $\mathbf{vExe}$  satisfies several conditions being the horizontal analogs of conditions for class  $\mathbf{uExe}$ ; Fig. 8 explains the idea. Hence, we have a category of data objects and *view mappings*,  $\mathbf{viewMap}$ , whose arrows are view execution tiles.

**B: Forward update translation.** Let  $\mathbf{w}$  be a  $\mathbf{wDef}$ -tile in Fig. 5(a). It provides a view definition and two update definitions, which can be executed for a given data instance  $A=(D_A, \tau_A)$ . There are two




Figure 8: View execution

concurrent computation path: We can get a pair of instances  $B$  and  $B'$  by executing view  $\mathbf{v}$  for  $A$  and for  $A' = A^{\mathbf{v}}$  respectively (see see Fig. 9(a)). Alternatively, we may execute update  $\mathbf{u}$  for view  $B = A^{\mathbf{v}}$  and get another instance  $B''$  Fig. 3(c).




Figure 9: Forward update translation

We call a tile  $\mathbf{w}$  *well-defined* if instances  $B'$  and  $B''$  are equal. Note that instance equality means both  $D_{B'} = D_{B''}$  and  $\tau_{B'} = \tau_{B''}$ . In the string-base notation, the law takes the form

$$(\text{Sync})_{\mathbf{w}} \quad (A^{\mathbf{v}})^{\mathbf{u}} = (A^{\mathbf{u}})^{\mathbf{p}}$$

where  $A \in [[\bullet \mathbf{w}]]$ ,  $\mathbf{p} = |\mathbf{w}$ ,  $\mathbf{u} = \mathbf{w}|$  and  $\overline{\mathbf{w}} = \underline{\mathbf{w}} = \mathbf{v}$ .

In this case, we may say that the view definition tile  $\mathbf{w}$  is executed for instance  $A$  and completes the pair  $(\mathbf{w}, A)$  up to an arrow *translation* cube shown in Fig. 9. We denote this cube by  $[\mathbf{w}]_A$ . Thus, a well-defined  $\mathbf{wDef}$ -tile gives rise to a correct translation cube.

We assume that every  $\mathbf{wDef}$ -tile is well-defined, and we thus have a class

$$\mathbf{wExe} = \{[\mathbf{w}]_A : \mathbf{w} \in \mathbf{wDef} \text{ and } A \in [[\bullet \mathbf{w}]]\}$$

of forward translation cubes: left face is translated into the right face.

**Definition 1** A well-defined *view system*  $\mathcal{V}$  is a graph  $\mathcal{U}$  with partial composition carrying the structure specified above: metadata constructs **dataDef**, **updDef**, **viewDef** and **wDef** (consisting of well-defined tiles), instance mapping assigning to each schema  $\mathcal{S}$  its set of instances  $[[\mathcal{S}]]$ , and execution mechanisms  $\mathbf{uExe}$ ,  $\mathbf{vExe}$  that result in data constructs **updMap**, **viewMap** and **wExe**; the latter defines an update translation execution mechanism.

### 3.3 Update policies

To set our main results about update policies, we will need the following categorical reformulation of the main definition.

**Proposition 1** *A well-defined view system gives rise to a double functor  $\tau: \mathbf{Inst} \rightarrow \mathbf{MData}$  between two thin double categories.*

*Proof (sketch).* We need to define two thin double categories and a functor between them. A thin double category is defined by its vertical and horizontal categories, and a class of tiles. For  $\mathbf{MData}$ , we take categories  $\mathbf{updDef}$  and  $\mathbf{viewDef}$  as  $\mathbf{MData}_1^v$  and  $\mathbf{MData}_1^h$  respectively. Double arrows are  $\mathbf{wDef}$ -tiles, and all the necessary special (idle and identity) tiles are there (see Fig. 5). The associativity of tile compositions immediately follows from associativity of composition in categories  $\mathbf{updDef}$  and  $\mathbf{viewDef}$ , and idle/identity tiles are neutral units because idle/identity arrows are such.

Definition of  $\mathbf{Inst}$  is a bit more intricate. 0-cells are instances,  $\mathbf{Inst}_0 \stackrel{\text{def}}{=} \bigcup_{S \in \mathbf{MData}_0} \llbracket S \rrbracket$ . Vertical arrows are update execution tiles,  $\mathbf{Inst}_1^v \stackrel{\text{def}}{=} \mathbf{uExe}$ , and they form a category by conditions in Section 3.2.4(A1). Horizontal arrows are view execution tiles,  $\mathbf{Inst}_1^h \stackrel{\text{def}}{=} \mathbf{vExe}$ , and they form a category by conditions in Section 3.2.4(A2). Double arrows are synchronous cubical diagrams from  $\mathbf{uExe}$ ,  $\mathbf{Inst}_2 \stackrel{\text{def}}{=} \mathbf{wExe}$ . Availability of all necessary idle/identity tiles (cubes), composition and its properties (associativity and units) are immediately checked because double arrows are thin (coincide with their cubical boundaries). Hence, we get a thin double category.

Typing double functor  $\tau$  is nothing but a projection that maps tiles constituting  $\mathbf{Inst}$ -arrows to their meta-sides, and maps cubes constituting  $\mathbf{Inst}$ -tiles to their meta-faces. Compatibility of this mapping with composition and identities is evident.  $\square$

The following result is crucial but is proved by immediate checking the definitions.

**Proposition 2** *The double functor  $\tau: \mathbf{Inst} \rightarrow \mathbf{MData}$  determined by a well-defined view system possesses the following inverting facilities (we write  $A: \mathcal{S}$  to say  $\tau(A) = \mathcal{S}$ ):*

(a) Update execution mechanism, which assigns to any vertical arrow (an update definition)  $\mathbf{u}: \mathcal{S} \rightarrow \mathcal{S}': \mathbf{MData}_1^v$  and data instance  $A: \mathcal{S}$ , an arrow  $[\mathbf{u}]_A: A: \mathcal{S} \rightarrow A': \mathcal{S}': \mathbf{Inst}_1^v$ ;

(b) View execution mechanism, which assigns to any horizontal arrow (view definition)

$\mathbf{v}: \mathcal{T} \rightarrow \mathcal{S}: \mathbf{MData}_1^h$  and instance  $A: \mathcal{S}$ , an arrow  $[\mathbf{v}]_A: B: \mathcal{T} \rightarrow A: \mathcal{S}: \mathbf{Inst}_1^h$ ;

(c) Translation execution mechanism, which assigns to any tile (translation definition)  $\mathbf{w} = (\mathbf{v}, \mathbf{u}, \mathbf{v}, \mathbf{p}) \in \mathbf{MData}_2$  and instance  $A: \bullet \mathbf{w}$ , a tile

$$[\mathbf{w}]_A = ([\mathbf{v}]_A, [\mathbf{u}]_{A^v}, [\mathbf{v}]_{A^u}, [\mathbf{p}]_A) \in \mathbf{Inst}_2.$$

Moreover, these execution operations are compatible with compositions:

$$(d) [x1; x2]_A = [x1]_B ; [x2]_A$$

where  $x1, x2$  are similar cells (either similar arrows or tiles), semicolon  $;$  is composition (vertical or horizontal), and  $B = A^{x2}$ . In categorical terms, it means that functors  $\tau_1^v$  and  $\tau_1^h$  are (split) opfibration and fibration respectively, and  $\tau$  is a double fibration [11].

**Corollary 3** *Let  $\mathbf{w}_{ij}$  ( $i = 1..m, j = 1..n$ ) be a net of  $\mathbf{wDef}$ -tiles like shown in Fig. 1. Let  $\mathbf{w} = \sum_{ij} \mathbf{w}_{ij}$  denotes the entire outer tile produced by composition of the inner tiles, and  $A \in \llbracket \bullet \mathbf{w} \rrbracket$ . Then  $[\mathbf{w}]_A = \sum_{ij} [\mathbf{w}_{ij}]_{A_{ij}}$ , where  $A_{ij}$  is the net of instances computed from  $A = A_{11}$  ( $A_{i+1,j} = A_i^{v_i}$ ,  $A_{i,j+1} = A_j^{u_j}$ ).*

Let  $\mathcal{V}$  be a well-defined view system,  $\mathbf{v} = \mathbf{v}_m; \mathbf{v}_{m-1}; \dots; \mathbf{v}_1: \mathcal{T} \rightarrow \mathcal{S}$  a complex view definition in  $\mathcal{V}$ , and  $\mathbf{U}_{\mathcal{T}}$  is the class of basic update definitions over  $\mathcal{T}$ . We remind that for any pair  $(\mathbf{v}, \mathbf{u}) \in \mathbf{viewDef}_1 \times \mathbf{updDef}_1$ , there is at least one  $\mathbf{wDef}$ -tile having  $\mathbf{u}$  as its right edge.

**Definition 2** *A backward translation policy for  $\mathcal{T}$  is a mapping  $\mathbf{T}: \mathbf{U}_{\mathcal{T}} \times \underline{m} \rightarrow \mathbf{wDef}$  with  $\underline{m} \stackrel{\text{def}}{=} \{1..m\}$ . A policy is called well-designed if it ensures tight tiling (see Fig. 1), i.e., formally,  $\mathbf{T}_{i+1}^u = |\mathbf{T}_i^u$  for all  $i=1, 2, \dots, m-1$  and  $\mathbf{T}_m^u = \mathbf{u}$ .*

**Proposition 4** *Let  $\mathbf{T}$  a well-designed update policy for view  $\mathbf{v}: \mathcal{T} \rightarrow \mathcal{S}$  as above, and  $\mathbf{u} = \mathbf{u}_1; \mathbf{u}_2; \dots; \mathbf{u}_n$ ,  $\mathbf{u}_j \in \mathbf{U}_{\mathcal{T}}$  a complex update over  $\mathcal{T}$ . Then for any instance  $A: \mathcal{S}$  and its view  $B = \mathbf{v}^A$ , the view update  $\bar{\mathbf{u}}_B: B \rightarrow B'$  is uniquely and correctly translated back to an update  $a: A \rightarrow A'$ .*

*Proof.* Decompositions  $\mathbf{v} = \mathbf{v}_m; \mathbf{v}_{m-1}; \dots; \mathbf{v}_1$  and  $\mathbf{u} = \mathbf{u}_1; \mathbf{u}_2; \dots; \mathbf{u}_n$  form a net, and the update policy “fills” each cell of this net with a translation definition  $\mathbf{wDef}$ -tile  $\mathbf{T}_i^j$ . Let  $\mathbf{w} = \sum_{ij} \mathbf{T}_i^j$  denotes the total composition of all these tiles,  $\mathbf{p} = |\mathbf{w}$  be its leftmost edge, and  $a = \bar{\mathbf{p}}_A: A \rightarrow A'$  is the corresponding update. We need to prove that translation is correct,

that is,  $\overline{v_{A'}} = B'$  or  $[\mathbf{w}]_{A'} = \overline{u_B}$ . We compute

$$\begin{aligned}
[\mathbf{w}]_{A'} &= [\sum_{ij} \mathbf{T}_i^j]_{A'} && \text{by definition of } \mathbf{w} \\
&= \sum_{ij} [\mathbf{T}_i^j]_{A_{ij}} && \text{by the corollary} \\
&= \sum_j \left( \sum_i [\mathbf{T}_i^j]_{A_{ij}} \right) && \text{by the interchange law} \\
&= \sum_j \left( [\sum_i \mathbf{T}_i^j]_{B_j} \right) && \text{by the corollary again} \\
&= \sum_j \overline{u_{j_{B_j}}} && \text{by definition of well} \\
&= \sum_j u_{j_{B_j}} && \text{designed policy} \\
&= \overline{u_B} && \text{by Proposition 2(d)} \\
&= \overline{u_B} && \text{by definition of } \mathbf{u}
\end{aligned}$$

Importantly, the proof shows that the result of backward translation does not depend on the way of composing tiles to obtain the total composition  $\mathbf{w}$ . This provides a space for optimization of update translations. The Pasting Lemma [13] allows us to extend the space for non-regular nets of view and update blocks.

## References

- [1] F. Bancilhon and N. Spyrtos. Update semantics of relational views. *TODS*, 6(4):557–575, 1981.
- [2] A. Bohannon, B. Pierce, and J. Vaughan. Relational lenses: a language for updatable views. In *PODS*, 2006.
- [3] R. Dawson and R. Pare. General associativity and general composition for double categories. *Cahiers de topologie et géométrie différentielle catégoriques*, 34:57–79, 1993.
- [4] U. Dayal and P. Bernstein. On the correct translation of update operations on relational views. *ACM TODS*, 7(3):381–416, 1982.
- [5] Z. Diskin. Algebraic models for bidirectional model synchronization. In *MoDELS*, pages 21–36, 2008.
- [6] Z. Diskin, K. Czarnecki, and M. Antkiewicz. Model-versioning-in-the-large: Algebraic foundations and the tile notation. In *ICSE 2009 Workshop on Comparison and Versioning of Software Models*, pages 7–12, 2009. DOI: 10.1109/CVSM.2009.5071715.
- [7] J. Foster, M. Greenwald, J. Moore, B. Pierce, and A. Schmitt. Combinators for bi-directional tree transformations: a linguistic approach to the view update problem. In *POPL*, pages 233–246, 2005.
- [8] J. N. Foster, M. Greenwald, C. Kirkegaard, B. Pierce, and A. Schmitt. Exploiting schemas in data synchronization. *J. Comput. Syst. Sci.*, 73(4):669–689, 2007.
- [9] J. N. Foster, M. Greenwald, J. Moore, B. Pierce, and A. Schmitt. Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem. *ACM Trans. Program. Lang. Syst.*, 29(3), 2007.
- [10] G. Gottlob, P. Paolini, and R. Zicari. Properties and update semantics of consistent views. *ACM TODS*, 13(4):486–524, 1988.
- [11] B. Jacobs. *Categorical logic and type theory*. Elsevier Science Publishers, 1999.
- [12] A. Keller. Algorithms for translating view updates to database updates for views involving selections, projections, and joins. In *PODS*, pages 154–163, 1985.
- [13] G.M. Kelly and R. Street. Review of the elements of 2-categories. In *Category Seminar, Sydney 1972/73*, Lecture Notes in Math., 420, pages 75–103, 1974.
- [14] R. Langerak. View updates in relational databases with an independent scheme. *ACM TODS*, 15(1):40–66, 1990.
- [15] E. Manes. *Algebraic Theories*. Graduate Text in Mathematics. Springer Verlag, 1976.
- [16] Anish Das Sarma, Omar Benjelloun, Alon Y. Halevy, Shubha U. Nabar, and Jennifer Widom. Representing uncertain data: models, properties, and algorithms. *VLDB J.*, 18(5):989–1019, 2009.
- [17] P. Stevens. Bidirectional model transformation in QVT: Semantic issues and open questions. In *MoDELS*, 2007.
- [18] Y. Xiong, D. Liu, Z. Hu, H. Zhao, M. Takeichi, and H. Mei. Towards automatic model synchronization from model transformations. In *ASE*, 2007.