



McMaster  
Centre for Software Certification

**Multiple Model Synchronization with Multiary  
Delta Lenses with Amendment and K-Output \***

Zinovy Diskin <sup>(✉)</sup> Harald König and Mark Lawford

# Multiple Model Synchronization with Multiary Delta Lenses with Amendment and K-Output \*\*

Zinovy Diskin<sup>1</sup> (✉) Harald König<sup>2</sup> and Mark Lawford<sup>1</sup>

<sup>1</sup> McMaster University, Hamilton, Canada

<sup>2</sup> University of Applied Sciences FHDW Hannover, Germany

diskinz@mcmaster.ca, harald.koenig@fhdw.de, lawford@mcmaster.ca

**Abstract.** Multiple (more than 2) model synchronization is ubiquitous and important for MDE, but its theoretical underpinning gained much less attention than the binary case. Specifically, the latter was extensively studied by the bx community in the framework of algebraic models for update propagation called *lenses*. Now we make a step to restore the balance and propose a notion of multiary delta lens. Besides multiarity, our lenses feature *reflective* updates, when consistency restoration requires some amendment of the update that violated consistency. We emphasize the importance of various ways of lens composition for practical applications of the framework, and prove several composition results.

## 1 Introduction

Modelling normally results in a set of inter-related models presenting different views of a single system at different stages of development. The former differentiation is usually referred to as “horizontal” (different views on the same abstraction level) and the latter as “vertical” (different abstraction levels beginning from the most general requirements down to design and further on to implementation). A typical modelling environment in a complex project is thus a collection of models (we will call them *local*) inter-related and inter-dependant along and across the horizontal and the vertical dimensions of the network. We will call the entire collection a *multimodel*, and refer to its component as to *local* models.

The system integrating local models can exist either materially (e.g., with UML modelling, a single UML model whose views are specified by UML diagrams, is physically stored by the UML tool) or virtually (e.g., several databases integrated into a federal database), or in a mixed way (e.g., in a complex modelling environment encompassing several UML models). Irrespective of the type of integration (material, virtual, mixed), the most fundamental property of a multimodel is its *global, or joint, consistency*: if local models do not contradict each other in their viewing of the system, then at least one system satisfying *all* local models exists; otherwise, we say local models are (globally) inconsistent.

---

\* Printed in *Formal Aspects of Computing* 31(5): 611-640 (2019) with multiple omissions in Sect. 7.1, which make it practically unreadable — blame Springer who sent the manuscript to production without final proofreading!

\*\* Printed in *Formal Aspects of Computing* 31(5): 611-640 (2019) with multiple omissions in Sect. 7.1, which make it practically unreadable — blame Springer who sent the manuscript to production without final proofreading!

If one of the local models changes and their joint consistency is violated, the related models should also be changed to restore consistency. This task of model synchronization is obviously of paramount importance for MDE, but its theoretical underpinning is inherently difficult and reliable automatic synchronization solutions are rare in practice. Much theoretical work partially supported by implementation has been done for the binary case (synchronizing two models) by the bidirectional transformation community (bx), specifically, by its TGG sub-community, see, e.g., [16]), and the *delta lens* sub-community on a more abstract level (delta lenses [11] can be seen as an abstract algebraic interface to TGG based synchronization [17]). However, disappointingly for practical applications, the case of multiary synchronization (the number of models to be synchronized is  $n > 2$ ) gained much less attention—cf. the energetic call to the community in a recent Stevens’ paper [32].

The context underlying bx is model transformation, in which one model in the pair is considered as a transform of the other even though updates are propagated in both directions (so called round-tripping). Once we go beyond  $n = 2$ , we switch to a more general context of inter-model relations beyond model-to-model transformations. Such situations have been studied in the context of multiview system consistency, see surveys [2,25], but rarely in the context of an accurate formal basis for update propagation. A notable exception is work by Trollmann and Albayrak [33,34,35]. In the first of these papers, they specify a grammar-based engine for generating consistent multimodels of arbitrary arity  $n \geq 2$ , with the case  $n = 2$  being managed by TGG and truly multiary cases  $n \geq 3$  are uniformly managed by what they call Graph-Diagram Grammars, GDG. In paper [34] they use GDG for building a multiary change propagation framework, which is close in its spirit to our framework developed in the paper but is much more concrete — we will provide a detailed comparison in the Related work section. Roughly, our framework of multiary delta lenses developed in the paper is to GDG-based update propagation as binary symmetric delta lenses are to TGG-based update propagation,  $\text{mxLens}/\text{GDG} \approx \text{bxLens}/\text{TGG}$ , where we refer to multiary update propagation as *mx* (contrasting it to binary bx). The latter relationship is described in [17]: binary delta lenses appear as an abstract algebraic interface to TGG-based change propagation; at some stage, we want to achieve similar results for mx-lenses and GDG (but not in this paper).

Our contributions to mx are as follows. We show with a simple example (Sect. 3) an important special feature of multiview modelling: consistency restoration may require not only update propagation to other models but the very update created inconsistency should itself be amended; thus, update propagation should, in general, be *reflective* (even for the case of a two-view system). Motivated by the example, in Sect. 4 we formally define the notion of a multi-model, and then in Sect. 5, give a formal definition of a *multiary* (symmetric) lens with amendment and state the basic algebraic laws such lenses must satisfy. Importantly, we have a special **KOutput** law that requires compatibility of update propagation with update composition for a restricted class  $\mathbf{K}^{\blacktriangleright}$  of composable update pairs.

Our major results are about lens composition. In Sect. 6, we define several operations over lenses, which produce complex lenses from simple ones: we first consider two forms of parallel composition in Sect. 6.1, and then two forms of sequential composition in Sections 6.2 and 6.4. Specifically, the construct of composing an  $n$ -tuple of asymmetric binary lenses sharing the same source

into a symmetric  $n$ -ary lens gives a solution to the problem of building mx synchronization via bx discussed by Stevens in [32].

We consider lens composition results crucially important for practical application of the framework. If a tool builder has implemented a library of elementary synchronization modules based on lenses and, hence, ensuring basic laws for change propagation, then a complex module assembled from elementary lenses will automatically be a lens and thus also enjoys the basic laws. This allows the developer to avoid additional integration testing, which can essentially reduce the cost of synchronization software.

The paper is an essential extension of our FASE'18 paper [7]. The main additions are i) a new section motivating our design choices, ii) a constrained Putput law (KPutput) and its thorough discussion, including a corresponding extension of the running example, iii) a counterexample showing that invertibility is not preserved by star composition, iv) two types of parallel composition of multiary lenses, v) Related Work and Future Work sections are essentially extended, particularly, an important subsection about multimodel updates including correspondence updates (categorification) is added.

## 2 Background: Design choices for the paper

In this section, we discuss our design choices for the paper: why we need multiarity, amendments, K-Putput, and why, although we recognize limitations of a framework only dealing with non-concurrent update scenarios, we still develop their accurate algebraic model in the paper.

**2.1 Why multiary lenses.** Consider, for simplicity, three models,  $A_1$ ,  $A_2$ , and  $A_3$ , working together (i.e., being models of the same integral system) and being in sync at some moment. Then one of the models, say,  $A_1$ , is updated to state  $A'_1$ , and consistency is violated. To restore consistency, the two other models are to be changed accordingly and we say that the update of model  $A_1$  is propagated to  $A_2$  and  $A_3$ . Thus, consistency restoration amounts to having three pairs of propagation operations,  $(\text{ppg}_{ij}, \text{ppg}_{ji})$ , with operation  $\text{ppg}_{ij}$  propagating updates of model  $A_i$  to model  $A_j$ ,  $i \neq j$ ,  $i, j = 1, 2, 3$ . It may seem that the synchronization problem can be managed by building three binary lenses  $\ell_{ij}$  – one lens per a pair of models.

However, when three models work together, their consistency is a ternary relation often irreducible to binary consistency relations. Figure 1 presents a simple example: three class diagrams shown in the figure are pairwise consistent while the whole triple is obviously inconsistent, i.e., violates a class diagram metamodel constraint (which prohibits composition cycles).

A binary lens can be seen as a couple of Mealy machines (we write  $A_1 \rightleftarrows A_2$ ) sharing a state space (say,  $R_{12}$ ). A ternary lens synchronizing a triple of models can also be seen as a triple of couples of Mealy machines ( $A_1 \rightleftarrows A_2$ ,  $A_1 \rightleftarrows A_3$ ,  $A_2 \rightleftarrows A_3$ ), but they share the same space  $R_{123}$  and hence mutually dependant on each other (they would be independent if each couple  $A_i \rightleftarrows A_j$  would have its own space  $R_{ij}$ ). Thus, multiple model synchronization is, in general, irreducible to chains of binary lenses and needs a new notion of a multiary lens.

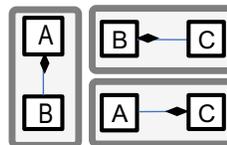


Fig. 1: Three jointly inconsistent class diagrams

**2.2 Why amendments.** Getting back to the example, suppose that the updated state  $A'_1$  goes beyond the projection of all jointly consistent states to the space  $A_1$ , and hence consistency cannot be restored with the first model being in state  $A'_1$ . When we work with two models, such cases could be a priori prohibited by modifying the corresponding metamodel  $M_1$  defining the model space so that state  $A'_1$  would violate  $M_1$ . When the number of models to sync grows, it seems more convenient and realistic to keep  $M_1$  more flexible and admitting  $A'_1$  but, instead, when the synchronizer is restoring consistency of all models, an amendment of state  $A'_1$  to a state  $A''_1$  is also allowed. Thus, update propagation works *reflectively* so that not only other models are changed but the initiating update from  $A_1$  to  $A'_1$  is itself amended and model  $A'_1$  is changed to  $A''_1$ . Moreover, in Sect. ?? we will consider examples of situations when if even state  $A'_1$  can be synchronized, a slight amendment still appears to be a better synchronization policy. For example, if consistency restoration with  $A'_1$  kept unchanged requires deletions in other models, while amending  $A'_1$  to  $A''_1$  allows to restore consistency by using additions only, then the update policy with amendments may be preferable – as a rule, additions are preferable to deletions.

Of course, allowing for amendments may open the Pandora box of pathological synchronization scenarios, e.g., we can restore consistency by rolling back the original update and setting  $A''_1 = A_1$ . We would like to exclude such solutions and bound amendments to work like completions of the updates rather than corrections and thus disallow any sort of “undoing”. To achieve this, we introduce a binary relation  $K \blacktriangleright$  of update compatibility: if an update  $u: A_1 \rightarrow A'_1$  is followed by update  $v: A'_1 \rightarrow A''_1$  and  $uK \blacktriangleright v$ , then  $v$  does not undo anything done by  $u$ . Then we require that the original update and its amendment be  $K \blacktriangleright$ -related. (Relation  $K \blacktriangleright$  and its formal properties are discussed in Sect. 4.2.)

**2.3 Why K-Putput.** An important and desired property of update propagation is its compatibility with update composition. If  $u$  and  $v$  are sequentially composable updates, and  $\text{put}$  is an update propagation operation, then its compositionality means  $\text{put}(u;v) = \text{put}(u); \text{put}(v)$ <sup>3</sup> —hence, the name Putput for the law. There are other equational laws imposed on propagation operations in the lens framework to guarantee desired synchronization properties and exclude unwanted scenarios. Amongst them, Putput is the most controversial: Putput without restrictions does not hold while finding an appropriate guarding condition – not too narrow to be practically usable and not too wide to ensure compositionality – has been elusive (cf. [15,18,10,20,6,5]). The practical importance of Putput follows from the possibilities of optimizing update propagation it opens: if Putput holds, instead of executing two propagations, the engine can execute just one. Moreover, before execution, the engine can optimize the procedure by preprocessing the composed update  $u;v$  and, if possible, converting it into an equivalent but easier manageable form (something like query optimization performed by database engines).

A preliminary idea of a constrained Putput is discussed in [20] under the name of a *monotonic* Putput: compositionality is only required for two consecutive deletions or two consecutive insertions (hence the term monotonic), which is obviously a too strong filter for typical practical applications. The idea of constraining Putput based on a compatibility relations over consecutive updates, e.g., relation  $K \blacktriangleright$  above, is much more flexible, and gives the name K-Putput

<sup>3</sup> to make this formula precise, some indexes are needed, but we have omitted them

for the law. It was proposed by Orejas *et al* in [29] for the binary case without amendment and intermodel correspondences (more accurately, with trivial correspondences being just pairs of models), and we adapt it for the case of full-fledge lenses with general correspondences and amendments. We consider our integration of  $\mathbf{K}\blacktriangleright$ -constrained Putput and amendments to be an important step towards making the lens formalism more usable and adaptable for practical tasks.

**2.4 Why non-concurrent synchronization.** In the paper, we will consider consistency violation caused by a change of only one model, and thus consistency is restored by propagating only one update, while in practice we often deal with several models changing concurrently. If these updates are independent, the case can be covered with one-update propagation framework using interleaving, but if concurrent updates are in conflict, consistency restoration needs a conflict resolution operation (based on some policy) and goes beyond the framework we will develop in the paper. One reason for this is technical difficulties of building lenses with concurrency – we need to specify reasonable equational laws regulating conflict resolution and its interaction with update composition. It would be a new stage in the development of the lens algebra.

Another reason is that the case of one-update propagation is still practically interesting and covers a broad class of scenarios – consider a UML model developed by a software engineer. Indeed, different UML diagrams are just different views of a single UML model maintained by the tool, and when the engineer changes one of the diagrams, the change is propagated to the model and then the changed model is projected to other diagrams. Our construct of star-composition of lenses (see Sect. 6.2 and Fig. 10) models exactly this scenario. Also, if a UML model is being developed by different teams concurrently, team members often agree about an interleaving discipline of making possibly conflicting changes, and the one-update framework is again useful. Finally, if concurrent changes are a priori known to be independent, this is well modelled by our construct of parallel composition of one-update propagating lenses (see Sect. 6.1).

**2.5 Lens terminology.** The domain of change propagation is inherently complicated and difficult to model formally. The lens framework that approaches the task is still under development and even the basic concepts are not entirely settled and co-exist in several versions (e.g., there are *strong* and *weak invertibility*, several versions of Putput, and different names for the same property of propagating idle updates to idle updates). This results in a diverse collection of different types of lenses, each of which has a subtype of so called *well-behaved* (*wb*) lenses (actually, several such as the notion of being *wb* varies), which are further branched into different notions of *very wb* lenses depending on the Putput version accepted. The diversity of lens types and their properties, on the one hand, and our goal to provide accurate formal statements, on the other hand, would lead to overly wordy formulations. To make them more compact, we use two bracket conventions.

Square brackets. If we say *A wb lens is called [weakly] invertible, when...*, we mean that using adjective *weakly* is optional for this paper as the only type of invertibility we consider is the weak invertibility. However, we need to mention it because the binary lens notion that we generalize in our multiary lens notion, is the weak invertibility to be distinguished from the strong one. Thus, we mention

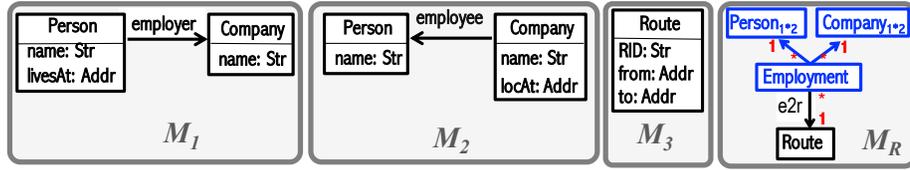


Fig. 2: Multi-Metamodel in UML

‘weakly’ in square brackets for the first time and then say just ‘invertible’ (except in a summarizing result like a theorem, in which we again mention [weakly]).

Round brackets. If a theorem reads *A span of (very) wb asymmetric lenses ... gives rise to a (very) wb symmetric lens ...* it means that actually we have two versions of the theorem: one is for *wb* lenses, and the other is for *very wb* lenses.

### 3 Example

We will consider a simple example motivating our framework. The formal constructs constituting the multiary delta lens framework will be illustrated with the example (or its fragments) and referred to as *Running example*. Although the lens framework is formal, the running example instantiating it, will be presented semi-formally: we will try to be precise enough, but an accurate formalization would require the machinery of graphs with diagram predicates and partial graph morphisms as described in our paper [26], and we do not want to overload this paper with formalities even more.

#### 3.1 A Multimodel to Play With

Suppose two data sources, whose schemas (we say metamodels) are shown in Fig. 2 as class diagrams  $M_1$  and  $M_2$  that record employment. The first source is interested in employment of people living in downtown, the second one is focused on software companies and their recently graduated employees. In general, population of classes **Person** and **Company** in the two sources can be different – they can even be disjoint, but if a recently graduated downtowner works for a software company, her appearance in both databases is very likely.

Now suppose there is an agency investigating traffic problems, which maintains its own data on commuting routes between addresses as shown by schema  $M_3$ . These data should be synchronized with commuting data provided by the first two sources and computable by an obvious relational join over  $M_1$  and  $M_2$ : roughly, the agency keeps traceability between the set of employment records and the corresponding commuting routes and requires their `from` and `to` attributes to be synchronized — below we will specify this condition in detail and explain the fourth metamodel  $M_R$  (see specification  $(R_{123})$  on the next page). In addition, the agency supervises consistency of the two sources and requires that if they both know a person  $p$  and a company  $c$ , then they must agree on the employment record  $(p, c)$ : it is either stored by both or by neither of the sources. For this synchronization, it is assumed that persons and companies are globally identified by their names. Thus, a triple of data sets (we will say models)  $A_1$ ,  $A_2$ ,  $A_3$ , instantiating the respective metamodels, can be either consistent (if the constraints described above are satisfied) or inconsistent (if they aren’t). In the

latter case, we normally want to change some or all models to restore consistency. We will call a collection of models to be kept in sync a *multimodel*.

To specify constraints for multimodels in an accurate way, we need an accurate notation. If  $A$  is a model instantiating metamodel  $M$  and  $X$  is a class in  $M$ , we write  $X^A$  for the set of objects instantiating  $X$  in  $A$ . Similarly, if  $r: X_1 \leftrightarrow X_2$  is an association in  $M$ , we write  $r^A$  for the corresponding binary relation over  $X_1^A \times X_2^A$ .<sup>4</sup> For example, Fig. 3 presents a simple model  $A_1$  instantiating  $M_1$  with  $\text{Person}^{A_1} = \{p_1, p'_1\}$ ,  $\text{Company}^{A_1} = \{c_1\}$ ,  $\text{employer}^{A_1} = \{(p_1, c_1)\}$ , and similarly for attributes, e.g.,

$$\text{livesAt}^{A_1} = \{(p_1, a1), (p'_1, a1)\} \subset \text{Person}^{A_1} \times \text{Addr}$$

( $\text{livesAt}^{A_1}$  and also  $\text{name}^{A_1}$  are assumed to be functions and  $\text{Addr}$  is the (model-independent) set of all possible addresses). Two other boxes present models  $A_2$  and  $A_3$  instantiating metamodels  $M_2$  and  $M_3$  resp.; we will discuss the rightmost box  $R$  later. The triple  $(A_1, A_2, A_3)$  is a (state of the) multimodel over the multimetamodel  $(M_1, M_2, M_3)$ , and we say it is *consistent* if the constraints specified below are satisfied.

Constraint (C1) specifies mutual consistency of models  $A_1$  and  $A_2$  in the sense described above:

$$(C1) \quad \begin{array}{l} \text{if } p \in \text{Person}^{A_1} \cap \text{Person}^{A_2} \text{ and } c \in \text{Company}^{A_1} \cap \text{Company}^{A_2} \\ \text{then } (p, c) \in \text{employer}^{A_1} \text{ iff } (c, p) \in \text{employee}^{A_2} \end{array}$$

Our other constraints specify consistency between the agency's data on commuting routes and the two data sources. We first assume a new piece of data that relates models: a relation  $(\text{e2r})^R$  whose domain is the integral set of employment records  $\text{Employment}^{A_1 * A_2}$  as specified below in  $(R_{123})$ :

$$(R_{123}) \quad \begin{array}{l} (\text{e2r})^R \subset \text{Employment}^{A_1 * A_2} \times \text{Route}^{A_3} \\ \text{where } \text{Employment}^{A_1 * A_2} \stackrel{\text{def}}{=} \text{employer}^{A_1} \cup (\text{employee}^{A_2})^{-1} \end{array}$$

This relation is described in the metamodel  $M_R$  in Fig. 2, where blue boxes denote (derived) classes whose instantiation is to be automatically computed (as specified above) rather than is given by the user.

To simplify presentation, we assume that all employees commute rather than work from home, which means that relation  $(\text{e2r})^R$  is left-total. Another simplifying assumption is that each employment record maps to exactly one commuting route, hence, relation  $(\text{e2r})^R$  is a single-valued mapping. Finally, several people living at the same address may work for the same company, which leads to different employment records mapped to the same route and thus injectivity is not required. Hence, we have the following intermodel constraint:

$$(C2) \quad \text{relation } (\text{e2r})^R \text{ is a total function } \text{Employment}^{A_1 * A_2} \longrightarrow \text{Route}^{A_3}$$

which is specified in the metamodel  $M_R$  in Fig. 2 by the corresponding multiplicities. However, the metamodel  $M_R$  as shown in Fig. 2, is still an incomplete specification of correspondences between models.

For each employment record  $e \in \text{Employment}^{A_1 * A_2}$ , there are defined the corresponding person

$$p = e.\text{employee}^{A_1 * A_2} \in \text{Person}^{A_1 * A_2} \stackrel{\text{def}}{=} \text{Person}^{A_1} \cup \text{Person}^{A_2}$$

<sup>4</sup> In general, an association  $r$  is interpreted as a multirelation  $r^A$ , but if the constraint [unique] is declared in the metamodel, then  $r^A$  must be a relation. We assume all associations in our metamodels are declared to be [unique] by default.

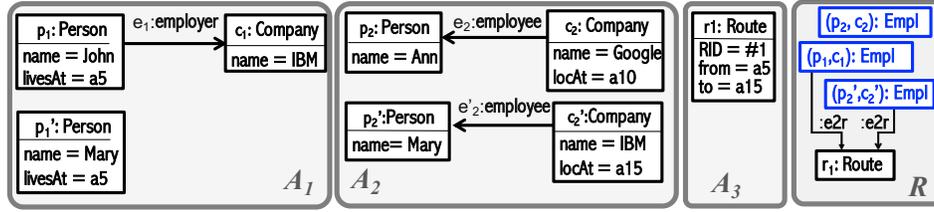


Fig. 3: Multimodel  $\mathcal{A}^\dagger = (A_1, A_2, A_3, R)$  with  $(p2p)^R = \{(p'_1, p'_2)\}$ ,  $(c2c)^R = \{(c_1, c'_2)\}$ , and  $(e2r)^R = \{(e_1, r_1), (e_2'^{-1}, r_1)\}$

with attribute

$$p.\text{livesAt}^{A_1 * A_2} = \begin{cases} p.\text{livesAt}^{A_1} & \text{if } p \in \text{Person}^{A_1}, \\ \perp & \text{otherwise.} \end{cases}$$

We thus assume that the domain  $\text{Addr}$  is extended with a bottom value/null  $\perp$ . Similarly, we define set  $\text{Company}^{A_1 * A_2}$  with attribute  $\text{locAt}^{A_1 * A_2}$  and the corresponding “end”  $e.\text{employer}^{A_1 * A_2} \in \text{Company}^{A_1 * A_2}$  for any employment record  $e$ . The latter is of special interest for the agency if both addresses,

$$e.\text{employee}^{A_1 * A_2}.\text{livesAt}^{A_1 * A_2} \text{ and } e.\text{employer}^{A_1 * A_2}.\text{locAt}^{A_1 * A_2},$$

are defined and thus define a certain commuting route to be consistent with its image  $e.(e2r)^R$  in  $\text{Route}^{A_3}$ .

More generally, the consistency of the two sets of commuting routes: that one derived from  $A_1$  and  $A_2$ , and that one stored in  $A_3$ , can be specified as follows:

$$(C3) \quad \begin{aligned} e.\text{employee}^{A_1 * A_2}.\text{livesAt}^{A_1 * A_2} &\leq e.(e2r)^R.\text{from}^{A_3} \\ e.\text{employer}^{A_1 * A_2}.\text{locAt}^{A_1 * A_2} &\leq e.(e2r)^R.\text{to}^{A_3} \end{aligned}$$

where inequality  $a_1 \leq a_2$  holds iff both values are certain and  $a_1 = a_2$ , or both values are nulls, or  $a_1$  is a null while  $a_2$  is certain.

Now it is easy to see that multimodel  $(A_1, A_2, A_3, R)$  in Fig. 3 is “two-times” inconsistent: (C1) is violated as both  $A_1$  and  $A_2$  know Mary and IBM, and  $(\text{IBM}, \text{Mary}) \in \text{employee}^{A_2}$  but  $(\text{Mary}, \text{IBM}) \notin \text{employer}^{A_1}$ , and (C2) is violated as  $A_2$  and  $R$  show an employment record  $e_2$  not mapped to a route in  $\text{Route}^{A_3}$ . Note also that if we map this record to the route #1 to fix (C2), then constraint (C3) will be violated. We will discuss consistency restoration in the next subsection, but first we need to finish our discussion of intermodel correspondences.

Note that correspondences between models have more data than specified in  $(R_{123})$ . Indeed, classes  $\text{Person}^{A_1}$  and  $\text{Person}^{A_2}$  are interrelated by a correspondence linking persons with the same name, and similarly for  $\text{Company}$  so that we have two partial injections:

$$(R_{12}) \quad (p2p)^R \subset \text{Person}^{A_1} \times \text{Person}^{A_2}, \quad (c2c)^R \subset \text{Company}^{A_1} \times \text{Company}^{A_2},$$

which are not shown in metamodel  $M_R$  (by purely technical reasons of keeping the figure compact and fitting in the page width). These correspondence links (we will write *corr-links*) may be implicit as they can always be restored using names as keys. In contrast, relation  $(e2r)$  is not determined by the component model states  $A_1, A_2, A_3$  and is an independent piece of data. Importantly, for given models  $A_{1,2,3}$ , there may be several different correspondence mappings  $(e2r)^R$

satisfying the constraints. For example, if there are several people living at the same address and working for the same company, all employment record can be mapped to the same route or to several different routes depending on how much carpooling is used. In fact, multiplicity of possible corr-specifications is a general story; it can happen for relations (p2p) and (c2c) as well if person and company names are not entirely reliable keys. Then we need a separate procedure of model matching or alignment that has to establish, e.g., whether objects  $p'_1 \in \text{Person}^{A_1}$  and  $p'_2 \in \text{Person}^{A_2}$  both named Mary represent the same real world object. Constraints we declared above implicitly involve corr-links, e.g., formula for (C1) is a syntactic sugar for the following formal statement: if  $(p_1, p_2) \in (\text{p2p})^R$  and  $(c_1, c_2) \in (\text{c2c})^R$  with  $p_i \in \text{Person}^{A_i}$ ,  $c_i \in \text{Company}^{A_i}$  ( $i = 1, 2$ ), then the following holds:  $(p_1, c_1) \in \text{employer}^{A_1}$  iff  $(c_2, p_2) \in \text{employee}^{A_2}$ . A precise formal account of this discussion can be found in [26].

Thus, a multimodel is actually a tuple  $\mathcal{A} = (A_1, A_2, A_3, R)$  where  $R$  is a correspondence specification (which, in our example, is a collection of correspondence relations  $(R_{123}), (R_{12})$  over sets involved). Consistency of a multimodel is a property of the entire 4-tuple  $\mathcal{A}$  rather than its 3-tuple carrier  $(A_1, A_2, A_3)$ .

### 3.2 Synchronization via Update Propagation

There are several ways to restore consistency of the multimodel in Fig. 3 w.r.t. constraint (C1). We may delete Mary from  $A_1$ , or delete her employment with IBM from  $A_2$ , or even delete IBM from  $A_2$ . We can also change Mary's employment from IBM to Google, which will restore (C1) as  $A_1$  does not know Google. Similarly, we can delete John's record from  $A_1$  and then Mary's employment with IBM in  $A_2$  would not violate (C1). As the number of constraints and the elements they involve increase, the number of consistency restoration variants grows fast.

The range of possibilities can be essentially decreased if we take into account the history of creating inconsistency and consider not only an inconsistent state  $\mathcal{A}^\dagger$  but update  $u: \mathcal{A} \rightarrow \mathcal{A}^\dagger$  that created it (assuming that  $\mathcal{A}$  is consistent). For example, suppose that initially model  $A_1$  contained record (Mary, IBM) (and  $A_3$  contained (a1, a15)-commute), and the inconsistency appears after Mary's employment with IBM was deleted in  $A_1$ . Then it's reasonable to restore consistency by deleting this employment record in  $A_2$  too; we say that deletion was propagated from  $A_1$  to  $A_2$ . If the inconsistency appears after adding (IBM, Mary)-employment to  $A_2$ , then it's reasonable to restore consistency by adding such a record to  $A_1$ . Although propagating deletions/additions to deletions/additions is typical, there are non-monotonic cases too. Let us assume that Mary and John are spouses and live at the same address, and that IBM follows an exotic policy prohibiting spouses to work together. Then we can interpret addition of (IBM, Mary)-record to  $A_2$  as swapping of the family member working for IBM, and then (John, IBM) is to be deleted from  $A_1$ .

Now let's consider how updates to and from model  $A_3$  may be propagated. As mentioned above, traceability/correspondence links play a crucial role here. If additions to  $A_1$  or  $A_2$  create a new commute, the latter has to be added to  $A_3$  (together with its corr-links) due to constraints (C2) and (C3). In contrast, if a new route is added to  $A_3$ , we may change nothing in  $A_{1,2}$  as (C2) does not require surjectivity of (e2r) (but further in the paper we will consider a more intricate policy). If a route is deleted from  $A_3$ , and it is traced via (e2r)<sup>R</sup> to one

or several corresponding employments in  $\text{Employment}^{A_1 * A_2}$ , then they are either deleted too, or perhaps remapped to other routes with the same from-topair of attributes if such exist. Similarly, deletions in  $\text{Employment}^{A_1 * A_2}$  may (but not necessarily) lead to the corresponding deletions in  $\text{Route}^{A_3}$  depending on the mapping  $(e2r)^R$ . Finally, updating addresses in  $A_1$  or  $A_2$  is propagated to the corresponding updates of from and to attributes in  $A_3$  to satisfy constraint (C3); similarly for attribute updates in  $A_3$ .

Clearly, many of the propagation policies above although formally correct, may contradict the real world changes and hence should be corrected, but this is a common problem of a majority of automatic synchronization approaches, which have to make guesses in order to resolve non-determinism inherent in consistency restoration.

### 3.3 Reflective Update Propagation

An important feature of update propagation scenarios above is that consistency could be restored without changing the model whose update caused inconsistency. However, this is not always desirable. Suppose again that violation of constraint (C1) in multimodel in Fig. 3 was caused by adding a new person Mary to  $A_1$ , e.g., as a result of Mary's moving to downtown. Now both models know both Mary and IBM, and thus either employment record (Mary, IBM) is to be added to  $A_1$ , or record (IBM, Mary) is to be removed from  $A_2$ . Either of the variants is possible, but in our context, adding (Mary, IBM) to  $A_1$  seems more likely and less specific than deletion (IBM, Mary) from  $A_2$ . Indeed, if Mary has just moved to downtown, the data source  $A_1$  simply may not have completed her record yet. Deletion (IBM, Mary) from  $A_2$  seems to be a different event unless there are strong causal dependencies between moving to downtown and working for IBM. Thus, an update policy that would keep  $A_2$  unchanged but amend addition of Mary to  $A_1$  with further automatic adding her employment for IBM (as per model  $A_2$ ) seems reasonable. This means that updates can be reflectively propagated (we also say self-propagated).

Of course, self-propagation does not necessarily mean non-propagation to other directions. Consider the following case: model  $A_1$  initially only contains (John, IBM) record and is consistent with  $A_2$  shown in Fig. 3. Then record (Mary, Google) was added to  $A_1$ , which thus became inconsistent with  $A_2$ . To restore consistency, (Mary, Google) is to be added to  $A_2$  (the update is propagated from  $A_1$  to  $A_2$ ) and (Mary, IBM) to be added to  $A_1$  as discussed above (i.e., addition of (Mary, Google) is both amended and propagated). Note, however, that in contrast to the previous case, now deletion of the record (IBM, Mary) from  $A_2$  looks like an equally reasonable scenario of Mary changing her employer. Thus, even for the simple case above, and the more complex cases of model interaction, the choice of the update policy (only amend, only propagate, or both) depends on the context, heuristics, and tuning the policy to practice.

A typical situation that needs an amendment facility is when the changes in interacting models have different granularity. With our simple running example, we can illustrate the point in the following (rather artificial) way. Suppose, again, that the record of Mary working for Google, and her address unknown (i.e.,  $\text{Mary.livesAt} = \perp$ ) is added to model  $A_1$ , and propagated to  $A_2$  as discussed above. Suppose that Google has a strict policy of only hiring those recent graduates who live on Bloor Street in Toronto downtown. Then in Mary's address

record, all fields could (and should!) be made certain besides the street number. Hence, adding Mary’s employment to model  $A_1$  should be amended with extending her address with data imposed by model  $A_2$ . For a more realistic example, consider model  $A_2$  specifying a complex engineering project  $P$  in the process of elaboration, while model  $A_1$  gives its very abstract view – the budget  $B$  of the project. If the budget changes from  $B$  to  $B'$ , the project should also be changed, but it is very likely that the budget of the changed project  $P'$  would be  $B''$  rather than exactly  $B'$ . A more general and formal description of this synchronization schema can be found in [5].

### 3.4 General schema

A general schema of update propagation including reflection is shown in Fig. 4. We begin with a consistent multimodel  $(A_1 \dots A_n, R)$ <sup>5</sup> one of which members is updated  $u_i: A_i \rightarrow A'_i$ . The propagation operation, based on a priori defined propagation policies as sketched above, produces:

- a) updates on all other models  $u'_j: A_j \rightarrow A''_j, 1 \leq j \neq i \leq n$ ;
- b) an amendment  $u_i^\circledast: A'_i \rightarrow A''_i$  to the original update;
- c) a new correspondence specification  $R''$  such that the updated multimodel

$$(A''_1 \dots A''_n, R'')$$

is consistent.

Below we introduce an algebraic model encompassing several operations and algebraic laws formally modelling situations considered so far.

## 4 Multimodel spaces

In this section we begin to build a formal framework for delta lenses: model spaces are described as categories whose objects are models and arrows are updates, which carry several additional relations and operations. We also abstractly define correspondences between models and our central notion of a (consistent) multimodel. We will follow an established terminological tradition (in the lens community) to give, first, a name to an algebra without any equational requirement, and then call an algebra satisfying certain equations *well-behaved*.

<sup>5</sup> Here we first abbreviate  $(A_1, \dots, A_n)$  by  $(A_1 \dots A_n)$ , and then write  $(A_1 \dots A_n, R)$  for  $((A_1 \dots A_n), R)$ . We will apply this style in other similar cases, and write, e.g.,  $i = 1 \dots n$  for  $i \in \{1, \dots, n\}$ .

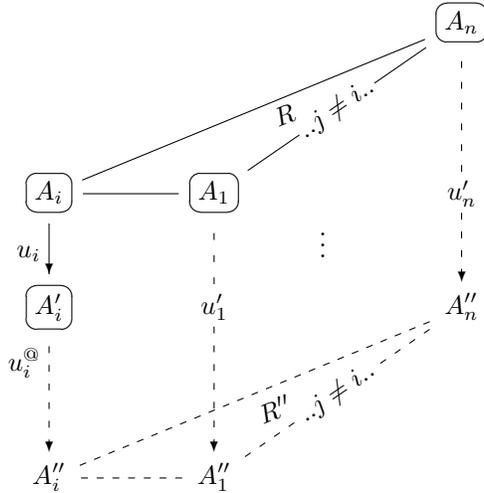


Fig. 4: Update propagation pattern (To distinguish given data from those produced by the operation, the former are shown with framed nodes and solid lines, while the latter are non-framed and dashed.)

#### 4.1 Background: Graphs, (co)Spans, and Categories

We reproduce well-known definitions to fix our notation. A (*directed multi-*)*graph*  $G$  consists of a set  $\text{Ob}(G)$  of *nodes* and a set  $\text{Ar}(G)$  of *arrows* equipped with two functions  $\mathbf{s}, \mathbf{t}: \text{Ar}(G) \rightarrow \text{Ob}(G)$  that give arrow  $a$  its *source*  $\mathbf{s}(a)$  and *target*  $\mathbf{t}(a)$  nodes. We write  $a: N \rightarrow N'$  if  $\mathbf{s}(a) = N$  and  $\mathbf{t}(a) = N'$ , and  $a: N \rightarrow \_$  or  $a: \_ \rightarrow N'$  if only one of these conditions is given.

Expressions  $G(N, N')$ ,  $G(N, \_)$ ,  $G(\_, N')$  denote sets of, resp., all arrows from  $N$  to  $N'$ , all arrows from  $N$ , and all arrows into  $N'$ .

A pair of arrows  $a_i: N \rightarrow N'_i$ ,  $i = 1, 2$ , with a common source is called a [*binary*] *span* with node  $N$  its *head* or *apex*, nodes  $N'_i$  *feet*, and arrows  $a_i$  *legs*. Dually, a pair of arrows  $a_i: N \leftarrow N'_i$ ,  $i = 1, 2$ , with a common target is called a (*binary*) *cospan* with apex, feet, and legs, defined similarly.

A [*small*] *category* is a graph, whose nodes are called *objects*, arrows are associatively composable, and every object has a special *identity* loop, which is the unit of the composition. In more detail, given two consecutive arrows  $a_1: \_ \rightarrow N$  and  $a_2: N \rightarrow \_$ , we denote the composed arrow by  $a_1; a_2$ . The identity loop of node  $N$  is denoted by  $\text{id}_N$ , and equations  $a_1; \text{id}_N = a_1$  and  $\text{id}_N; a_2 = a_2$  are to hold. We will denote categories by bold letters, say,  $\mathbf{A}$ , and often write  $A \in \mathbf{A}$  rather than  $A \in \text{Ob}(\mathbf{A})$  for its objects.

A *functor* is a mapping of nodes and arrows from one category to another, which respects sources and targets as well as identities and composition. Having a tuple/family of categories  $\mathbf{A} = (\mathbf{A}_1 \dots \mathbf{A}_n)$ , their *product* is a category  $\prod \mathbf{A} = \mathbf{A}_1 \times \dots \times \mathbf{A}_n$  whose objects are tuples  $A = (A_1 \dots A_n) \in \text{Ob}(\mathbf{A}_1) \times \dots \times \text{Ob}(\mathbf{A}_n)$ , and arrows from  $(A_1 \dots A_n)$  to  $(A'_1 \dots A'_n)$  are tuples of arrows  $u = (u_1 \dots u_n)$  with  $u_i: A_i \rightarrow A'_i$  for all  $i = 1 \dots n$ .

#### 4.2 Model Spaces and updates

Basically, a *model space* is a category, whose nodes are called *model states* or just *models*, and arrows are (*directed*) *deltas* or *updates*. For an arrow  $u: A \rightarrow A'$ , we treat  $A$  as the state of the model before update  $u$ ,  $A'$  as the state after the update, and  $u$  as an update specification. Structurally, it is a specification of correspondences between  $A$  and  $A'$ . Operationally, it is an edit sequence (edit log) that changed  $A$  to  $A'$ . The formalism does not prescribe what updates are, but assumes that they form a category, i.e., there may be different updates from state  $A$  to state  $A'$ ; updates are composable; and idle updates  $\text{id}_A: A \rightarrow A$  (doing nothing) are the units of the composition. A prominent example of model spaces is the category of graphs where updates are encoded as (certain equivalence classes of) binary spans between them. They are heavily used in the theory of Graph Transformations [13]. In this way an update  $u: A \rightarrow A'$  can be a deletion or an addition or a combination of both.

We require every model space  $\mathbf{A}$  to be endowed with two additional constructs of update compatibility.

**Sequential compatibility of updates** We assume a family  $(\mathbf{K}_A^{\blacktriangleright})_{A \in \text{Ob}(\mathbf{A})}$  of binary relations  $\mathbf{K}_A^{\blacktriangleright} \subset \mathbf{A}(\_, A) \times \mathbf{A}(A, \_)$  indexed by objects of  $\mathbf{A}$ , and specifying *non-conflicting* or *compatible* consecutive updates. Intuitively, an update  $u$  into  $A$  is compatible with update  $u'$  from  $A$ , if  $u'$  does not revert/undo anything done by  $u$ , e.g., it does not delete/create objects created/deleted by  $u$ , or



of *sequential compatibility* relations for sequential update pairs as described above. The third and fourth component are tightly coupled:

$$\mathbf{K}^\blacktriangle = (\mathbf{K}_A^\blacktriangle \subset \mathbf{A}(A, \_) \times \mathbf{A}(A, \_) \mid A \in \text{Ob}(\mathbf{A}))$$

is a family of *concurrently compatible* or *mergeable* spans of updates, and

$$\text{merge} = (\text{merge}_A \mid A \in \text{Ob}(\mathbf{A}))$$

is a family of merge operations as shown in (1). The domain of operation  $\text{merge}_A$  is exactly the set  $\mathbf{K}_A^\blacktriangle$ . We will denote the components of merge by  $\text{merge}_{A,i}$  with  $i = 1, 2$  and omit index  $A$  if they are clear from the context or not important. Writing  $\text{merge}_A(u_1, u_2)$  implicitly assumes  $(u_1, u_2) \in \mathbf{K}_A^\blacktriangle$ .  $\square$

**Definition 2 (Well-behaved Model Spaces).**

- ( $\mathbf{K}^\blacktriangleright$  and Id) For all  $u \in \mathbf{A}(\_, A)$ ,  $u' \in \mathbf{A}(A, \_)$ :  $(u, \text{id}_A), (\text{id}_A, u') \in \mathbf{K}_A^\blacktriangleright$ ,  
 For any three consecutive updates  $u, v, w$ , we require:
- ( $\mathbf{K}^\blacktriangleright \mathbf{K}^\blacktriangleright$ )<sub>1</sub>  $(u \mathbf{K}^\blacktriangleright v \wedge v \mathbf{K}^\blacktriangleright w) \text{ imply } (u \mathbf{K}^\blacktriangleright v \wedge uv \mathbf{K}^\blacktriangleright w)$ ,
- ( $\mathbf{K}^\blacktriangleright \mathbf{K}^\blacktriangleright$ )<sub>2</sub>  $(uv \mathbf{K}^\blacktriangleright w \wedge u \mathbf{K}^\blacktriangleright v) \text{ imply } (v \mathbf{K}^\blacktriangleright w \wedge u \mathbf{K}^\blacktriangleright vw)$ ,
- where composition is denoted by concatenation ( $uv$  for  $u; v$  etc).
- ( $\mathbf{K}^\blacktriangle$  and  $\mathbf{K}^\blacktriangleright$ ) For all  $(u_1, u_2) \in \mathbf{K}_A^\blacktriangle$ :  $(u_i, \tilde{u}_i) \in \mathbf{K}_{A'_i}^\blacktriangleright$ ,  $i = 1, 2$ , where  $\tilde{u}_i = \text{merge}_{A,i}(u_1, u_2)$
- (MergeSym) For all  $u_1: A \rightarrow A'_1$ ,  $u_2: A \rightarrow A'_2$ , if  $(\tilde{u}_1, \tilde{u}_2) = \text{merge}_A(u_1, u_2)$ ,  
 then  $(\tilde{u}_2, \tilde{u}_1) = \text{merge}_A(u_2, u_1)$
- (Mergeld) For all  $u: A \rightarrow A'$ :  $\text{merge}_A(u, \text{id}_A) = (\text{id}_{A'}, u)$

$\square$

The first condition is the already discussed natural property for sequential compatibility. The pair of conditions below it requires complementing updates not to revert anything done by the other update. The last three conditions are obvious requirements to the operation merge.

We assume all our model spaces to be wb and, as a rule, will omit explicit mentioning. Each of the metamodels  $M_1..M_3$  in the running example gives rise to a (wb) model space of its instances as discussed above.

### 4.3 Correspondences and Multimodels

We will work with families of model spaces indexed by a finite set  $I$ , whose elements can be seen as space *names*. To simplify notation, we will assume that  $I = \{1, \dots, n\}$  although ordering will not play any role in our framework. Given a tuple of model spaces  $\mathbf{A}_1, \dots, \mathbf{A}_n$ , we will refer to objects and arrows of the product category  $\prod \mathbf{A} = \mathbf{A}_1 \times \dots \times \mathbf{A}_n$  as *tuple models* and *tuple updates*, and denote them by letters without indexes, e.g.,  $u: A \rightarrow A'$  is a family  $u_i: A_i \rightarrow A'_i$ ,  $i = 1..n$ . We will call components of tuple models and updates their *feet*. We also call elements of a particular space  $\mathbf{A}_i$  *foot models* and *foot updates*.

**Definition 3 (Multispaces, Alignment, Consistency).** Let  $n \geq 2$  be a natural number. An *n-ary multi-space* is a triple  $\mathcal{A} = (\mathbf{A}, \text{Corr}, \partial)$  with the following components. The first one is a tuple of (wb) model spaces  $\mathbf{A} = (\mathbf{A}_1, \dots, \mathbf{A}_n)$  called the *boundary* of  $\mathcal{A}$ . The other two components is a class **Corr** of elements

called (*consistent*) *correspondences* or *corrs*, and a family of *boundary mappings*  $\partial = (\partial_i: \text{Corr} \rightarrow \text{Ob}(\mathbf{A}_i) \mid i = 1..n)$ .

A corr  $R \in \text{Corr}$  is understood as a correspondence specification interrelating models  $A_i = \partial_i R$ ; the latter are also called *R's feet*, and we say that models  $A_i$  are aligned via  $R$ . We write  $\partial R$  for the tuple  $(\partial_1 R \dots \partial_n R)$ . When we consider several multispaces and need an explicit reference, we will write  $\partial \mathcal{A}$  for the boundary  $\mathbf{A}$  of the entire multispace, and  $\text{Corr}(\mathcal{A})$  for its class of corrs.

Given a model tuple  $A = (A_1 \dots A_n) \in \text{Ob}(\prod \mathbf{A})$ , we write  $\text{Corr}(A)$  for the set  $\{R \in \text{Corr} \mid \partial R = A\}$ , and  $\text{Corr}(A_i)$  for the class  $\{R \in \text{Corr} \mid \partial_i R = A_i\}$ ; thus,  $\text{Corr}(A) = \bigcap_{i=1..n} \text{Corr}(A_i)$ .<sup>6</sup>

Although in this paper all corrs are considered consistent by default, we will often mention their consistency explicitly to recall the context and to ease comparison with similar but a bit more general frameworks, in which inconsistent corrs are also considered [11,6].  $\square$

**Definition 4 (Multimodels).** A (*consistent*) *multimodel* over a multispace  $\mathcal{A}$  is a couple  $\mathcal{A} = (A, R)$  of a model tuple  $A = (A_1 \dots A_n)$  with a (consistent) corr  $R \in \text{Corr}(A)$  relating the component models. A *multimodel update*  $\mathbf{u}: \mathcal{A} \rightarrow \mathcal{A}'$  is a pair  $(u, (R, R'))$ , whose first component is a tuple update  $u = (u_i: \partial_i R \rightarrow \partial_i R' \mid i = 1..n)$ , and the second component is a pair of the old and the new corrs. Identity updates are pairs  $(\text{id}, (R, R))$ , whose tuple component consists of identities  $\text{id}_{\partial_i R}$ ,  $i = 1..n$ , only. It is easy to check that so defined multimodels and their updates determine a category that we denote by  $\mathcal{R}$ .  $\square$

*Remark 1.* The notions of a corr and a multimodel are actually synonyms: any corr  $R$  is simultaneously a multimodel  $(\partial R, R)$ , and any multimodel  $(A, R)$  is basically just a corr  $R$  as the feet part of the notion is uniquely restored by setting  $A = \partial R$ , thus,  $\text{Corr} \cong \text{Ob}(\mathcal{R})$ . We can extend the equivalence to arrows too by defining corr updates exactly as we defined multimodel updates. This would make class  $\text{Corr}$  into a category  $\mathbf{Corr}$  isomorphic to  $\mathcal{R}$ . Using two symbols and two names for basically the same notion is, perhaps, confusing but we decided to keep them to keep track of the historical use of the terminology. The choice of the word depends on the context: if we focus on the  $R$  component of pair  $(A, R)$ , we say ‘‘corr’’, if we focus on the  $A$  component, we say ‘‘multimodel’’.

Thus, for this paper, multimodel updates are basically tuple updates  $u$  while their corr-component  $(R, R')$  is trivial (co-discrete in the categorical jargon) and so we actually will not need the category  $\mathbf{Corr}$  (or  $\mathcal{R}$ ) explicitly declared. However, the notion will be useful when we will discuss categorification of the framework in Future Work Sect. 7.1, and in Related Work Sect. 8.

*Example 1.* The *Running example* of Sect.3 gives rise to a 3-ary multimodel space. For  $i = 1..3$ , space  $\mathbf{A}_i$  consists of all models instantiating metamodel  $M_i$  in Fig.2 and their updates. Given a model tuple  $A = (A_1, A_2, A_3)$ , a consistent corr  $R \in \text{Corr}(A)$  is given by a triple of relations  $(\text{p2p})^R$ ,  $(\text{c2c})^R$ , and  $(\text{e2r})^R$  (the first is specified by formula  $(R_{123})$  on p. 7 and the other two by  $(R_{12})$  on p. 8) such that the intermodel constraints (C1-3) are satisfied. If we rely on person and company names as keys, then relations  $(\text{p2p})^R$ ,  $(\text{c2c})^R$  are derived from foot models  $A_1$  and  $A_2$ , but if these keys are not reliable, then the two relations are independent components of the multimodel. Relation  $(\text{e2r})^R$  is always an independent component.

<sup>6</sup> With this line of notation, the entire class  $\text{Corr}(\mathcal{A})$  could be denoted by  $\text{Corr}(\mathbf{A})$ .

## 5 Update Propagation and Multiary (Delta) Lenses

Update policies described in Sect. 3 can be extended to cover propagation of all updates  $u_i$ ,  $i \in 1..3$  according to the pattern in Fig. 4. This is a non-trivial task, but after it is accomplished, we obtain a synchronization framework, which we algebraically model as an algebraic structure called a (*very*) *well-behaved lens*. In this term, *lens* refers to a collection of diagram operations defined over a multispace of models, each of which takes a configuration of models, corrs and updates, and returns another configuration of models, corrs and updates — then we say that the operation *propagates* updates. A lens is called *well-behaved*, if its propagation operations satisfy a set of algebraic laws specified by equations. This terminological discipline goes back to the first papers, in which lenses were introduced [14]. We define and discuss well-behaved lenses in the next Sect. 5.1. Additionally, in Sect. 5.2, we discuss yet another important requirement to a reasonable synchronization framework: compatibility of update propagation with update composition, which is specified by the most controversial amongst the lens laws – the (in) famous Putput. We define a suitably constrained version of the law and call it Kputput, and call a well-behaved lens satisfying the KPutput law *very well-behaved* (again following the terminological tradition of [14]).

### 5.1 Well-behaved lenses

**Definition 5 (Symmetric lenses).** An  $n$ -ary *symmetric lens* is a pair  $\ell = (\mathcal{A}, \text{ppg})$  with  $\mathcal{A}$  an  $n$ -ary multimodel called the *carrier* of  $\ell$ , and  $\text{ppg} = (\text{ppg}_i \mid i = 1..n)$  a family of operations of the following arities. Operation  $\text{ppg}_i$  takes a consistent corr  $R$  with boundary  $\partial R = (A_1 \dots A_n)$ , and a foot update  $u_i: A_i \rightarrow A'_i$  as its input, and returns three data items (a,b,c) specified below.

(a) an  $(n - 1)$ -tuple of updates  $u'_j: A_j \rightarrow A''_j$  with  $1 \leq j \neq i \leq n$ ;

(b) an *amendment*  $u_i^\circledast: A'_i \rightarrow A''_i$  to the original update so that we have a new *reflective* update  $u'_i = u_i$ ;  $u_i^\circledast: A_i \rightarrow A''_i$

These data define a tuple update  $u': A \rightarrow A'' = (A''_1, \dots, A''_i, \dots, A''_n)$ .

(c) a new corr  $R' \in \text{Corr}(A'')$ .

In fact, operation  $\text{ppg}_i$  completes a (local) foot update  $u_i$  to a (global) update of the entire multimodel  $\mathbf{u}: \mathcal{A} \rightarrow \mathcal{A}'$ , whose components are  $(u'_j)_{j \neq i}$ ,  $u'_i = u_i$ ;  $u_i^\circledast$ , and the pair  $(R, R')$  (see also Fig. 4).  $\square$

Note that all  $\text{ppg}$  operations are only defined for consistent corrs and return consistent corrs. The latter requirement is often formulated as a special lens law (often called *Correctness*) but in our framework, it is embedded in the arity of propagation operations.

**Notation.** If the first argument  $R$  of operation  $\text{ppg}_i$  is fixed, the corresponding family of unary operations (whose only argument is  $u_i$ ) will be denoted by  $\text{ppg}_i^R$ . By taking the  $j$ th component of the multi-element result, we obtain single-valued unary operations  $\text{ppg}_{ij}^R$  producing, resp. updates  $u'_j = \text{ppg}_{ij}^R(u_i): A_j \rightarrow A''_j$  for all  $j \neq i$  (see clause (a) of the definition) while  $\text{ppg}_{ii}^R$  returns the amendment  $u_i^\circledast$ . We also have operation  $\text{ppg}_{i\star}^R$  returning a new consistent corr  $R'' = \text{ppg}_{i\star}^R(u_i)$  according to (c).

**Definition 6 (Closed updates).** Given a lens  $\ell = (\mathcal{A}, \text{ppg})$  and a corr  $R \in \text{Corr}(A_1 \dots A_n)$ , we call an update  $u_i: A_i \rightarrow A'_i$  *R-closed*, if  $\text{ppg}_{ii}^R(u_i) = u_i$  (i.e.,  $u_i^\circledast = \text{id}_{A'_i}$ ). An update  $u_i$  is *closed* if it is *R-closed* for all  $R \in \text{Corr}(A_i)$ . Lens  $\ell$  is called *closed* at foot  $\mathbf{A}_i$ , if all updates in  $\mathbf{A}_i$  are  $\ell$ -closed.  $\square$

**Definition 7 (Well-behaved lenses).** A lens  $\ell = (\mathcal{A}, \text{ppg})$  is called *well-behaved* (*wb*) if the following laws hold for all  $i = 1..n$ ,  $A_i \in \text{Ob}(\mathbf{A}_i)$ ,  $R \in \text{Corr}(A_i)$  and  $u_i: A_i \rightarrow A'_i$ , cf. Fig. 4

$$(\text{Stability})_i \quad \text{ppg}_{ij}^R(\text{id}_{A_i}) = \text{id}_{A_j} \text{ for all } j = 1..n, \text{ and } \text{ppg}_{i*}^R(\text{id}_{A_i}) = R$$

$$(\text{Reflect1})_i \quad (u_i, u_i^\circledast) \in \mathbf{K}_{A'_i}^{\blacktriangleright}$$

$$(\text{Reflect2})_i \quad \text{ppg}_{ij}^R(u_i; u_i^\circledast) = \text{ppg}_{ij}^R(u_i) \text{ for all } j \neq i$$

$$(\text{Reflect3})_i \quad \text{ppg}_{ii}^R(u_i; u_i^\circledast) = \text{id}_{A'_i} \text{ where } A'_i \text{ is the target model of } u_i^\circledast,$$

where in laws (Reflect1–3),  $u_i^\circledast$  stands for  $\text{ppg}_{ii}^R(u_i)$

Stability says that lenses do nothing voluntarily. Reflect1 says that amendment works towards “completion” rather than “undoing”, and Reflect2-3 are idempotency conditions to ensure the completion indeed done.

**Definition 8 (Invertibility).** A wb lens is called [*weakly*] *invertible*, if it satisfies the following law for any  $i$ , update  $u_i: A_i \rightarrow A'_i$  and  $R \in \text{corr}A_i$ :

$$(\text{Invert})_i \quad \text{for all } j \neq i: \text{ppg}_{ij}^R(\text{ppg}_{ji}^R(\text{ppg}_{ij}^R(u_i))) = \text{ppg}_{ij}^R(u_i) \quad \square$$

This law deals with “round-tripping”: operation  $\text{ppg}_{ji}^R$  applied to update  $u_j = \text{ppg}_{ij}^R(u_i)$  results in update  $\hat{u}_i$  equivalent to  $u_i$  in the sense that  $\text{ppg}_{ij}^R(\hat{u}_i) = \text{ppg}_{ij}^R(u_i)$  (see [11] for a motivating discussion).

*Example 2 (Trivial lenses  $\mathfrak{t}_n(\mathbf{A})$ ).* A category consisting of one object and one (necessary identity) arrow is called *terminal*. All terminal categories are isomorphic; we fix one, whose object is denoted by  $\mathbf{1}$  while the category is denoted by  $\mathbf{1}$  (bold 1). The terminal category  $\mathbf{1}$  gives rise to a unique *terminal space* with  $\mathbf{K}_1^{\blacktriangleright} = \{\text{id}_1, \text{id}_1\}$ ,  $\mathbf{K}_1^\blacktriangle = \{\text{id}_1, \text{id}_1\}$  and  $\text{merge}(\text{id}_1, \text{id}_1) = \text{id}_1$ .

Any model space  $\mathbf{A}$  gives rise to the following *trivial  $n$ -ary lens*  $\mathfrak{t}_n(\mathbf{A})$ . The first foot space  $\mathbf{A}_1 = \mathbf{A}$  while for all  $j = 2..n$ ,  $\mathbf{A}_j = \mathbf{1}$ . Tuple models are uniquely determined by their first foot, and given such a model  $(A, 1..1)$ , the set of corrs is the singleton  $\{A\}$  with  $\partial_1 A = A$  and  $\partial_{2..n} A = \mathbf{1}$ , and this only corr is considered consistent. Hence, all mutlimodels are consistent and update propagation is not actually needed. For any  $\mathbf{A}$ , lens  $\mathfrak{t}_n(\mathbf{A})$  is a wb, invertible lens closed at all its feet in a trivial way.  $\square$

The next example is more interesting.

*Example 3 (Identity Lenses  $\text{id}_n(\mathbf{A})$ ).* Let  $\mathbf{A}$  be an arbitrary model space. It generates an  $n$ -ary lens  $\text{id}_n(\mathbf{A})$  as follows. The carrier  $\mathcal{A}$  has  $n$  identical feet spaces:  $\partial_i \mathcal{A} = \mathbf{A}$  for all  $i = 1..n$ . The corr set  $\text{Corr}A$  for a tuple model  $A = (A_1..A_n)$  is the singleton  $\{A\}$  with  $\partial_i A = A_i$ ; this corr is consistent iff  $A_1 = A_2 = \dots = A_n$ . All updates are propagated to themselves (hence the name *identity lens*). Obviously,  $\text{id}_n(\mathbf{A})$  is a wb, invertible lens closed at all its feet.  $\square$

## 5.2 Very well-behaved lenses

We consider an important property of update propagation—its compatibility with update composition. A simple compositionality law would require that the composition  $u_i; v_i$  of two consecutive foot updates  $u_i: A_i \rightarrow A'_i$ ,  $v_i: A'_i \rightarrow B_i$ , is propagated into the composition of propagations:

$$\text{ppg}_{ij}^R(u_i; v_i) = \text{ppg}_{ij}^R(u_i); \text{ppg}_{ij}^{R'}(v_i)$$

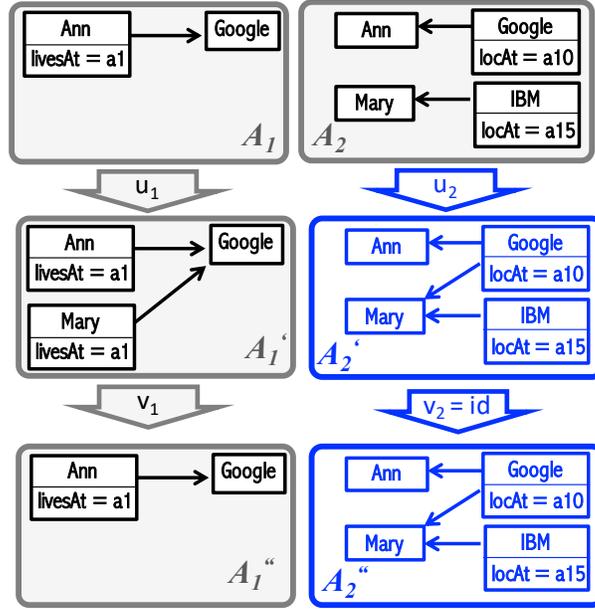


Fig. 6: Unconstrained Putput is violated - Example 1

with  $R''$  being the corr provided by the first propagation,  $R'' = \text{ppg}_{i^*}^R(u_i)$ . It is however well known that such a simple law (called PutPut) often does not hold.

Figure 6 presents a simple example (we use a more compact notation, in which values of the attribute name are used as OIDs – the primary key attribute idea). At the initial moment, the binary multimodel  $(A_1, A_2)$  with the (implicit) corr given by name matching is consistent: the only intermodel constraint (C1) is satisfied. Update  $u_1$  adds a new employment record to model  $A_1$ , constraint (C1) is violated, and to restore consistency, a new employment is added to  $A_2$  by update  $u_2 = \text{ppg}_{12}^R(u_1)$  (the new propagated model and update are shown with blue lines and blank). Then update  $v_1$  deletes the record added by  $u_1$ , but as the resulting multimodel  $(A_1', A_2')$  remains consistent, nothing should be done (if we follow the Hippocraticness principle in bx introduced by Stevens [31]) and thus  $v_1$  is propagated to identity,  $v_2 = \text{ppg}_{12}^{R''}(v_1) = \text{id}_{A_2'}$ . Now we notice that the composition  $(u_1; v_1)$  is identity, and hence is to be propagated to identity, i.e.,  $\text{ppg}_{12}^R(u_1; v_1) = \text{id}_{A_2}$ , while  $\text{ppg}_{12}^R(u_1); \text{ppg}_{12}^{R''}(v_1) = u_2; v_2 = u_2 \neq \text{id}_{A_2}$ , and  $A_1'' \neq A_2''$ .<sup>7</sup>

Obviously, the violation story will still hold if models  $A_1$  and  $A_2$  are much bigger and updates  $u_1$  and  $v_1$  are parts of bigger updates (but Mary should not appear in  $A_1$ ).

Figure 7 presents a more interesting ternary example. Multi-model  $\mathcal{A}$  is consistent, but update  $u_3$  deletes route #1 and violates constraint (C2). In this case,

<sup>7</sup> In more detail, equality of models  $A_1$  and  $A_1''$  is a bit more complicated than shown in the figure. Objects Ann in  $A_1$  and Ann in  $A_1''$  will actually have different OIDs, but as OIDs are normally invisible, we consider models up to their isomorphism w.r.t. OID permutations that keep attribute values unchanged. Hence, models  $A_1$  and  $A_1''$  are isomorphic and become equal after we factorize models by the equivalence described above.

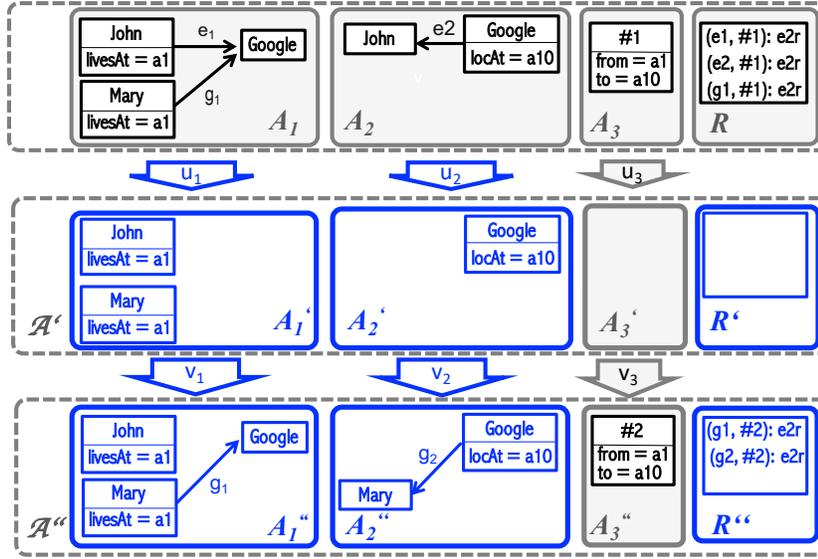


Fig. 7: Unconstrained Putput is violated—Example 2

the most direct propagation policy would delete all employment records related to #1 but, to keep changes minimal, would keep people in  $A_1$  and companies in  $A_2$  as specified by model  $\mathcal{A}'$  in the figure. Then a new route is inserted by update  $v_3$  (note the new OID), which has the same from and to attributes as route #1. There are multiple ways to propagate such an insertion. The simplest one is to do nothing as the multimodel ( $A_1', A_2', A_3', R' = \emptyset$ ) is consistent (recall that mapping (e2r) is not necessarily surjective). However, we can assume a more intelligent (non-Hippocratic) policy that looks for people with address a1 in the first database (and finds John and Mary), then checks somehow (e.g., by other attributes in the database schema  $M_1$  which we did not show) if some of them are recent graduates in IT (and, for example, finds that Mary is such) and thus could work for a company at the address a10 (Google). Then such a policy would propagate update  $v_3$  to  $v_1$  and  $v_2$  as shown in the figure. Now we notice that the composition of updates  $u_3$  and  $v_3$  is an identity<sup>8</sup>, and thus  $u_3; v_3$  would be propagated to the identity on the multimodel  $\mathcal{A}$ . However,  $\mathcal{A}'' \neq \mathcal{A}$  and Putput fails. (Note that it would also fail if we used a simpler policy of propagating update  $v_3$  to identity updates on  $A_1'$  and  $A_2'$ .)

A common feature of the two examples is that the second update  $v$  fully reverts the effect of the first one,  $u$ , while their propagations do not fully enjoy such a property due to some “side effects”. In fact, Putput is violated if update  $v$  would just partially revert update  $u$ . Several other examples and a more general discussion of this phenomenon that forces Putput to fail can be found in [5]. However, it makes sense to require compositional update propagation for two *sequentially compatible* updates in the sense of Def. 1; we call the respective version of the law KPutput law with ‘K’ recalling the sequential compatibility relation  $\mathbf{K} \blacktriangleright$ . To manage update amendments, we will also need to use update merging operation as shown in Fig. 8

<sup>8</sup> see the previous footnote 7

**Definition 9 (Very well-behaved lenses).** A wb lens is called *very well behaved* (*vwb*), if it satisfies the following KPutput law for any  $i$ ,  $\text{corr } R \in \text{corr } A_i$ , and updates  $u_i: A_i \rightarrow A'_i$ ,  $v_i: A'_i \rightarrow B_i$  (see Fig. 8).

Let  $u_i^\circledast = \text{ppg}_{ii}^R(u_i)$ . Suppose that  $(u_i, v_i) \in \mathbf{K}_{A'_i}^{\blacktriangleright}$  and  $(v_i, u_i^\circledast) \in \mathbf{K}_{A'_i}^{\blacktriangleleft}$ , then the merge  $(\tilde{v}_i, \tilde{u}_i^\circledast) = \text{merge}(v_i, u_i^\circledast)$  is defined (see Fig. 8) and the following equalities are required to hold:

$$\begin{aligned} (\text{KPutput})_{j \neq i} \quad & \text{ppg}_{ij}^R(u_i; v_i) = \text{ppg}_{ij}^R(u_i); \text{ppg}_{ij}^{R''}(\tilde{v}_i), \\ & \text{and } (\text{ppg}_{ij}^R(u_i), \text{ppg}_{ij}^{R''}(\tilde{v}_i)) \in \mathbf{K}_{A_j''}^{\blacktriangleright}, \end{aligned}$$

$$(\text{KPutput})_{ii} \quad (u_i; v_i)^\circledast = \tilde{u}_i^\circledast; (\tilde{v}_i)^\circledast$$

see the dashed arrows in Fig. 8, which depict the propagation of the composed updates.  $\square$

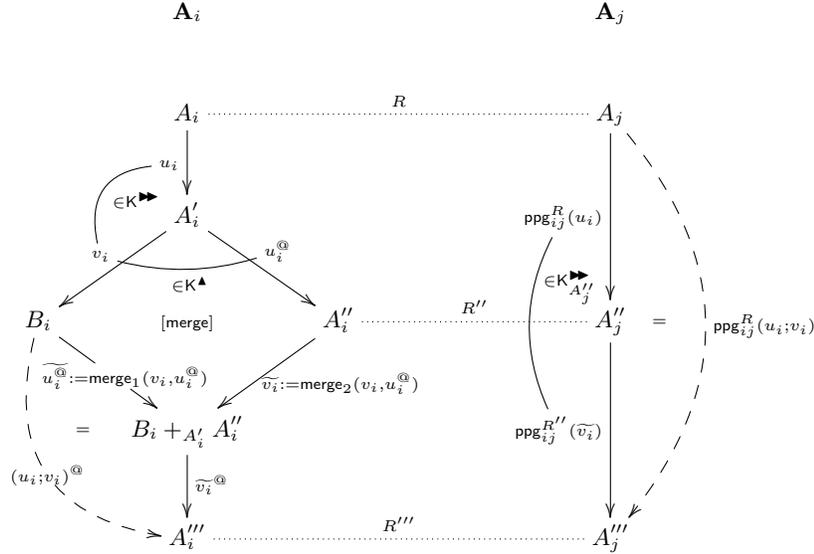


Fig. 8: (KPutput)-Law (for update  $x$ , expression  $x^\circledast$  stands for the amendment  $\text{ppg}_{ii}^{R(x)}(x)$  with  $R(x)$  being the corr at the source of  $x$  – it labels the corresponding horizontal dotted line)

**Corollary 1 (Closed vwb lenses).** For a vwb lens as defined above, if update  $u_i$  is  $R$ -closed (i.e.,  $u_i^\circledast = \text{id}_{A'}$ ), then the following equations hold:

$$\begin{aligned} (\text{KPutput})_{j \neq i}^{\text{closed}} \quad & \text{ppg}_{ij}^R(u_i; v_i) = \text{ppg}_{ij}^R(u_i); \text{ppg}_{ij}^{R''}(v_i), \\ & \text{and } (\text{ppg}_{ij}^R(u_i), \text{ppg}_{ij}^{R''}(v_i)) \in \mathbf{K}_{A_j''}^{\blacktriangleright}, \end{aligned}$$

$$(\text{KPutput})_{ii}^{\text{closed}} \quad (u_i; v_i)^\circledast = v_i^\circledast$$

## 6 Compositionality of Model Synchronization: Playing Lego with Lenses

We study how lenses can be composed. Assembling well-behaved and well-tested small components into a bigger one is a cornerstone of software engineering. If a mathematical theorem guarantees that desired properties carry over from the components to the composition, additional integration tests for checking these properties are no longer necessary. This makes lens composition results practically important.

In Sect. 6.1, we consider parallel composition of lenses, which is easily manageable. Sequential composition, in which different lenses share some of their feet, and updates propagated by one lens are taken and propagated further by one or several other lenses, is much more challenging and considered in the other two subsections. In Sect. 6.2, we consider "star-composition", cf. Fig. 10 and show that under certain additional assumptions (very) well-behavedness carries over from the components to the composition. However, invertibility does not carry over to the composition – we show this with a counterexample in Sect. 6.3. In Sect. 6.4, we study how (symmetric) lenses can be assembled from asymmetric ones and prove two easy theorems on the property preservation for such composition.

Since we now work with several lenses, we need a notation for lens' components. Given a lens  $\ell = (\mathcal{A}, \text{ppg})$ , we write  $\mathbf{A}^\ell$ ,  $\mathbf{A}(\ell)$  or  $\partial\ell$  for  $\partial\mathcal{A}$ ,  $\text{Corr}^\ell$  or  $\text{Corr}(\ell)$  for  $\text{Corr}(\mathcal{A})$ , and  $\partial_i^\ell(R)$  for the  $i$ -th boundary of  $\text{corr } R$ . Propagation operations of the lens  $\ell$  are denoted by  $\ell.\text{ppg}_{ij}^R$ ,  $\ell.\text{ppg}_{i\star}^R$ . We will often identify an aligned multimodel  $(A, R)$  and its  $\text{corr } R$  as they are mutually derivable (see Remark 1 on p.15).

We will also need the notion of lens isomorphism.

**Definition 10 (Isomorphic lenses).** Two  $n$ -ary lenses  $\ell$  and  $\ell'$  are *isomorphic*,  $\ell \cong \ell'$ , if

- (a) their feet are isomorphic via a family of isomorphism functors  $f_i: \mathbf{A}_i^\ell \rightarrow \mathbf{A}_i^{\ell'}$ ;
- (b) their classes of corrs are isomorphic via bijection  $f_{\text{Corr}}: \text{Corr}^\ell \rightarrow \text{Corr}^{\ell'}$  commuting with boundaries:  $\partial_i^{\ell'}(f_{\text{Corr}}(R)) = f_i(\partial_i^\ell(R))$  for all  $i$ ;
- (c) their propagation operations are compatible with isomorphisms above: for any foot update  $u_i: A_i^\ell \rightarrow B_i^{\ell'}$  and  $\text{corr } R \in \text{Corr}^\ell(A_i)$  for lens  $\ell$ , we have

$$f_j(\ell.\text{ppg}_{ij}^R(u_i)) = \ell'.\text{ppg}_{ij}^{f_{\text{Corr}}(R)}(f_i(u_i)).$$

That is, two composed mappings, one propagates  $u_i$  with lens  $\ell$  then maps the result to  $\ell'$ -space, and the other maps  $u_i$  to  $\ell'$ -space and then propagates it with lens  $\ell'$ , produce the same result.

### 6.1 Parallel Lego: Lenses working in parallel

We will consider two types of parallel composition. The first is chaotic (co-discrete in the categorical parlance). Suppose we have several clusters of synchronized models, i.e., models within the same cluster are synchronized but models in different clusters are independent. We can model such situations by considering several lenses  $\ell_1, \dots, \ell_k$ , each one working over its own multimodel space  $\mathcal{A}_i$ . Although mutually independent w.r.t data, clusters are time-related

and it makes sense to talk about multimodels  $\mathcal{A}_1^t, \dots, \mathcal{A}_k^t$  coexisting at some time moment  $t$ : such a tuple of multimodels can be seen as the state of the multimodel  $\mathcal{A}_1 \times \dots \times \mathcal{A}_k$  at moment  $t$ . As model clusters are data-independent, we can propagate tuples of updates  $(u_1, \dots, u_k)$  — one update per cluster, to other such tuples. For example, if we have a ternary lens  $\kappa$  with feet  $\mathbf{A}_i, i = 1..3$ , and a binary lens  $\ell$  with feet  $\mathbf{B}_j, j = 1, 2$ , any pair of updates  $(u_i, v_j) \in \mathbf{A}_i \times \mathbf{B}_j$  can be propagated to pairs  $(u_{i'}, v_{j'}) \in \mathbf{A}_{i'} \times \mathbf{B}_{j'}$  with  $i' = 1..3$  and  $j' = 1, 2$  by the two lenses working in parallel: lens  $\kappa$  propagates  $u_i$  and  $\ell$  propagates  $v_j$ . The resulting synchronization can be seen as a six-ary lens with feet  $\mathbf{A}_i \times \mathbf{B}_j$ . A general construction that composes lenses  $\ell_1, \dots, \ell_k$  of arities  $n_1, \dots, n_k$  resp., into a product lens  $\ell_1 \times \dots \times \ell_k$  of arity  $n_1 \times n_2 \times \dots \times n_k$  is described in Sect. 6.1.1

Our second construct of parallel composition is for lenses of the same arity working in a strongly coordinated way. Suppose that our traffic agency has several branches in different cities, all structured in a similar way, i.e., over metamodels  $M_i, i = 1, 2, 3$  in Fig. 2. However, now we have families of models  $A_i^x$  with  $x$  ranging over cities. Suppose also a strong discipline of coordinated updates, in which all models of the same type, i.e., with a fixed metamodel index  $i$  but different city index  $x$ , are updated simultaneously. Then global updates are tuples like  $(u_1^x: A_1^x \rightarrow A_1'^x \mid x \in \text{Cities})$  or  $(u_2^x: A_2^x \rightarrow A_2'^x \mid x \in \text{Cities})$ . Such tuples can be propagated componentwise, i.e., city-wise, so that we have a global ternary lens, whose each foot is indexed by cities. Thus, in contrast to the chaotic parallel composition, the arity of the coordinated composed lens equals to the arity of components. We will formally define the construct in Sect. 6.1.2.

### 6.1.1. Chaotic Parallel Composition.

**Definition 11.** Let  $\kappa$  and  $\ell$  be two lenses of arities  $m$  and  $n$ . We first choose the following two-dimensional enumeration of their product  $mn$ : any number  $1 \leq i \leq mn$  is assigned with two natural numbers as specified below:

$$(2) \quad i \mapsto (i_\kappa, i_\ell) = \begin{cases} (1, i) & \text{if } 1 \leq i \leq n, \\ (2, i) & \text{if } n < i \leq 2n, \\ \dots & \\ (m, i) & \text{if } (m-1)n < i \leq mn, \end{cases}$$

Of course, we could choose another such enumeration but its only effect is reindexing/renaming the feet while synchronization as such is not affected. Now we define the *chaotic parallel composition* of  $\kappa$  and  $\ell$  as the  $m \times n$ -ary lens  $\kappa \times \ell$  with Boundary spaces:  $\partial_i^{\kappa \times \ell} = (\partial_{i_\kappa}^\kappa, \partial_{i_\ell}^\ell)$

Corrs:  $\text{Corr}(\kappa \times \ell) = \text{Corr}(\kappa) \times \text{Corr}(\ell)$  with boundaries  
 $\partial_i^{(\kappa \times \ell)}(Q, R) = (\partial_{i_\kappa}^\kappa(Q), \partial_{i_\ell}^\ell(R))$  for all  $i$

Operations: Given an update  $\mathbf{u}_i$  at foot  $i$  of lens  $\kappa \times \ell$ , i.e., a pair of updates  $(u_{i_\kappa}, v_{i_\ell})$  with  $u_{i_\kappa}: A_{i_\kappa} \rightarrow A'_{i_\kappa}, v_{i_\ell}: B_{i_\ell} \rightarrow B'_{i_\ell}$ , and corrs  $Q \in \text{Corr}^\kappa(A_{i_\kappa}), R \in \text{Corr}^\ell(B_{i_\ell})$ , we define  
 $(\kappa \times \ell). \text{ppg}_{ij}^{(Q,R)}(\mathbf{u}_i) \stackrel{\text{def}}{=} (\kappa. \text{ppg}_{i_\kappa j_\kappa}^Q(u_{i_\kappa}), \ell. \text{ppg}_{i_\ell j_\ell}^R(v_{i_\ell}))$   
 $(\kappa \times \ell). \text{ppg}_{i_*}^{(Q,R)}(\mathbf{u}_i) \stackrel{\text{def}}{=} (\kappa. \text{ppg}_{i_\kappa *}^Q(u_{i_\kappa}), \ell. \text{ppg}_{i_\ell *}^Q(v_{i_\ell})).$

Furthermore, for any models  $A \in \text{Ob}(\mathbf{A})$  and  $B \in \text{Ob}(\mathbf{B})$ , relations  $\mathbf{K}_{A \times B}^{\blacktriangleright\blacktriangleright}$  and  $\mathbf{K}_{A \times B}^\blacktriangle$  are the obvious rearrangement of elements of  $\mathbf{K}_A^\blacktriangleright \times \mathbf{K}_B^\blacktriangleright$  and  $\mathbf{K}_A^\blacktriangle \times \mathbf{K}_B^\blacktriangle$ .  $\square$

**Lemma 1.** *If  $\kappa$  and  $\ell$  are (very) wb (and invertible), then  $\kappa \times \ell$  is (very) wb (and invertible).*

*Proof.* All verifications can be carried out componentwise.  $\square$

**Lemma 2.** *Chaotic parallel composition is associative up to isomorphism:  $(\kappa \times \ell) \times \ell' \cong \kappa \times (\ell \times \ell')$*

*Proof.* Straightforward based on associativity of the Cartesian product.  $\square$

The two lemmas imply

**Theorem 1 (Chaotic Parallel Composition).** *Let  $\ell_1, \dots, \ell_k$  be a tuple of lenses of arities  $n_i$ ,  $i = 1..k$ . Then  $n_1 \times \dots \times n_k$ -ary lens  $\ell_1 \times \dots \times \ell_k$  is defined [up to isomorphism], and it is (very) wb (and invertible) as soon as all lenses  $\ell_i$  are such.*  $\square$

### 6.1.2. Coordinated Parallel Composition.

**Definition 12.** Let  $\kappa$  and  $\ell$  be two lenses of the same arity  $n$ . Their *coordinated parallel composition* is the  $n$ -ary lens  $\kappa || \ell$  with

- Boundary spaces:  $\partial_i^{\kappa || \ell} = (\partial_i^\kappa, \partial_i^\ell)$  for all  $1 \leq i \leq n$ ,
- Corrs:  $\text{Corr}(\kappa || \ell) = \text{Corr}(\kappa) \times \text{Corr}(\ell)$  with boundaries  $\partial_i^{\kappa || \ell} = (\partial_i^\kappa, \partial_i^\ell)$  for all  $i$
- Operations: If  $\mathbf{u}_i = (u_i: A_i \rightarrow A'_i, v_i: B_i \rightarrow B'_i)$  is an update at the  $i$ -th foot  $(A_i, B_i)$  of lens  $\kappa || \ell$ , and  $Q \in \text{Corr}^\kappa(A_i)$ ,  $R \in \text{Corr}^\ell(B_i)$  are corrs, then

$$(\kappa || \ell). \text{ppg}_{ij}^{(Q,R)}(\mathbf{u}_i) \stackrel{\text{def}}{=} (\kappa. \text{ppg}_{ij}^Q(u_i), \ell. \text{ppg}_{ij}^R(v_i))$$

and

$$(\kappa || \ell). \text{ppg}_{i\star}^{(Q,R)}(\mathbf{u}_i) \stackrel{\text{def}}{=} (\kappa. \text{ppg}_{i\star}^Q(u_i), \ell. \text{ppg}_{i\star}^R(v_i)).$$

Furthermore, for any models  $A \in \text{Ob}(\mathbf{A})$  and  $B \in \text{Ob}(\mathbf{B})$ , relations  $\mathbf{K}_{A \times B}^{\blacktriangleright}$  and  $\mathbf{K}_{A \times B}^{\blacktriangle}$  are the obvious rearrangement of elements of  $\mathbf{K}_A^{\blacktriangleright} \times \mathbf{K}_B^{\blacktriangleright}$  and  $\mathbf{K}_A^{\blacktriangle} \times \mathbf{K}_B^{\blacktriangle}$ .  $\square$

**Lemma 3.** *If  $\kappa$  and  $\ell$  are (very) wb (and invertible), then  $\kappa || \ell$  is (very) wb (and invertible).*

*Proof.* All verifications can be carried out componentwise.  $\square$

**Lemma 4.** *Coordinated parallel composition is associative up to isomorphism:  $(\kappa || \ell) || \ell' \cong \kappa || (\ell || \ell')$*

*Proof.* Straightforward based on associativity of the Cartesian product.  $\square$

The two lemmas imply

**Theorem 2 (Coordinated Parallel Composition).** *Let  $\ell_1, \dots, \ell_k$  be a tuple of lenses of the same arity  $n$ . Then  $n$ -ary lens  $\ell_1 || \dots || \ell_k$  is defined [up to isomorphism], and it is (very) wb (and invertible) as soon as all lenses  $\ell_i$  are such.*  $\square$

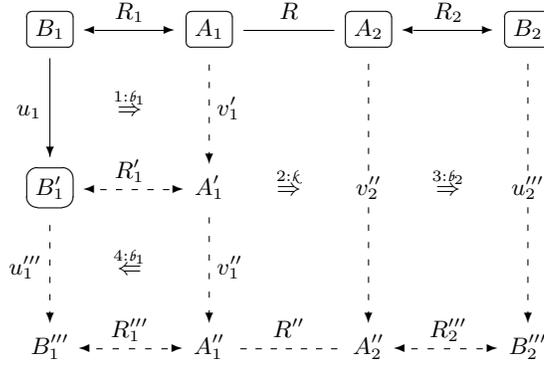


Fig. 9: Running example via lenses

## 6.2 Sequential Lego 1: Star Composition

**Running Example Continued.** Diagram in Fig. 9 presents a refinement of our example, which explicitly includes relational storage models  $B_{1,2}$  for the two data sources. We assume that object models  $A_{1,2}$  are simple projective views of databases  $B_{1,2}$ : data in  $A_i$  are copied from  $B_i$  without any transformation, while additional tables and attributes that  $B_i$ -data may have are excluded from the view  $A_i$ . Synchronisation of bases  $B_i$  and their views  $A_i$  can be realized by simple *constant-complement* lenses  $b_i$ ,  $i = 1, 2$  (see, e.g., [24]), such that consistent corrs  $R_i \in \text{Corr}(b_i)(B_i, A_i)$  (in fact traceability mappings) are exactly those for which the projection of  $B_i$  yields  $A_i$ .

Finally, let  $\kappa$  be a lens synchronizing models  $A_1, A_2, A_3$  as described in Sect. 3, and  $R \in \text{Corr}(\kappa)(A_1, A_2, A_3)$  be a corr for some  $A_3$  not shown in the figure.

Consider the following update propagation scenario. Suppose that at some moment we have consistency  $(R_1, R, R_2)$  of all five models, models  $A_{1,2,3}$  are as shown in Fig. 3 except that model  $A_1$  (and the base  $B_1$ ) do not have any data about Mary. Then model  $B_1$  is updated with  $u_1: B_1 \rightarrow B'_1$  that, say, adds to  $B_1$  a record of Mary working for Google. Consistency is restored with a four-step propagation procedure shown by double-arrows labeled by  $x:y$  with  $x$  the step number and  $y$  the lens doing the propagation.

**Step 1:** lens  $b_1$  propagates update  $u_1$  to  $v'_1$  that adds (Mary, Google) to view  $A_1$  with no amendment to  $u_1$  as  $v'_1$  is just a projection of  $u_1$ , thus,  $B'_1 = B''_1$ . Note also the updated traceability mapping  $R'_1 \in \text{Corr}(b_1)(B'_1, A'_1)$

**Step 2:** lens  $\kappa$  propagates  $v'_1$  to  $v''_2$  that adds (Google, Mary) to  $A_2$ , and amends  $v'_1$  with  $v''_1$  that adds (Mary, IBM) to  $A'_1$  to satisfy constraint (C1); a new consistent corr  $R''$  is also computed.

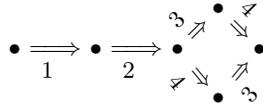
**Step 3:** lens  $b_2$  propagates  $v''_2$  to  $u''_2$  that adds Mary's employment by Google to  $B_2$  with, perhaps, some other specific relational storage changes not visible in  $A_2$ . We assume no amendment to  $v''_2$  as otherwise access to relational storage would amend application data. Thus we have a consistent corr  $R'''_2$  as shown.

**Step 4:** lens  $b_1$  maps update  $v''_1$  (see above in Step 2) backward to  $u''_1$  that adds (Mary, IBM) to  $B'_1$  so that  $B'''_1$  includes both (Mary, Google) and (Mary,

IBM) and a respective consistent corr  $R_1'''$  is provided. There is no amendment for  $v_1''$  by the same reason as in Step 3.

Thus, all five models in the bottom line of Fig. 9 ( $A_3''$  is not shown) are mutually consistent and all show that Mary is employed by IBM and Google. Synchronization is restored, and we can consider the entire scenario as propagation of  $u_1$  to  $u_2'''$  and its amendment with  $u_1'''$  so that finally we have a consistent corr  $(R_1''', R'', R_2''')$  interrelating  $B_1''', A_3'', B_2'''$ . Amendment  $u_1'''$  is compatible with  $u_1$  as nothing is undone and condition  $(u_1, u_1''') \in K_{B_1}^{\blacktriangleright}$  holds; the other two equations required by Reflect2-3 for the pair  $(u_1, u_1''')$  also hold. For our simple projection views, these conditions will hold for other updates too, and we have a well-behaved propagation from  $B_1$  to  $B_2$  (and trivially to  $A_3$ ). Similarly, we have a wb propagation from  $B_2$  to  $B_1$  and  $A_3$ . Propagation from  $A_3$  to  $B_{1,2}$  is non-reflective and done in two steps: first lens  $\kappa$  works, then lenses  $b_i$  work as described above (and updates produced by  $\kappa$  are  $b_i$ -closed). Thus, we have built a wb ternary lens synchronizing spaces  $\mathbf{B}_1, \mathbf{B}_2$  and  $\mathbf{A}_3$  by joining lenses  $b_1$  and  $b_2$  to the central lens  $\kappa$ .

**Discussion.** Reflection is a crucial aspect of lens composition. The diagram below describes the scenario above as a transition system and shows that Steps 3 and 4 can be performed in either order.



It is the non-trivial amendment created in Step 2 that causes the necessity of Step 4, otherwise Step 3 would finish consistency restoration (with Step 4 being an idle transition). On the other hand, if update  $v_2''$  in Fig. 9 would not be closed for lens  $b_2$ , we would have yet another concurrent step complicating the scenario. Fortunately for our example with simple projective views, Step 4 is simple and provides a non-conflicting amendment, but the case of more complex views beyond the constant-complement class needs care and investigation. Below we specify a simple situation of lens composition with reflection a priori excluded, and leave more complex cases for future work.

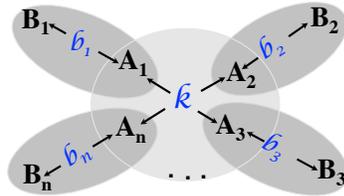


Fig. 10: Star Composition

**Formal Definition of Star Composition** Suppose we have an  $n$ -ary lens  $\kappa = (\mathcal{A}, \kappa.\text{ppg})$  with  $\mathcal{A}$  based on the model space tuple  $\mathbf{A} = (\mathbf{A}_1 \dots \mathbf{A}_n)$ , and for every  $i \leq n$  a binary lens  $b_i = (\mathbf{B}_i, \mathbf{A}_i, b_i.\text{ppg})$ , with the second model space

$\mathbf{A}_i$  being the  $i$ th model space of  $\kappa$  (see Fig.10, where  $\kappa$  is depicted in the center and  $b_i$  are shown as ellipses adjoint to  $\kappa$ 's feet). We also assume that the following *Junction Conditions* holds for all  $i \leq n$ :

(Junction) <sub>$i$</sub>  All updates propagated to  $\mathbf{A}_i$  by lens  $b_i$  are  $\kappa$ -closed,  
and all updates propagated to  $\mathbf{A}_i$  by lens  $\kappa$  are  $b_i$ -closed.

Below we will write the sextuple of operations  $b_i.\text{ppg}^{R_i}$  as the family

$$(b_i.\text{ppg}_{xy}^{R_i} \mid x \in \{\mathbf{A}, \mathbf{B}\}, y \in \{\mathbf{A}, \mathbf{B}, \star\}).$$

Likewise we write  $\partial_x^{b_i}$  with  $x \in \{\mathbf{A}, \mathbf{B}\}$  for the boundary functions of lenses  $b_i$ .

The above configuration gives rise to the following  $n$ -ary lens  $\ell$ . The carrier is the tuple of model spaces  $\mathbf{B}_1 \dots \mathbf{B}_n$  together with their already contained compatible consecutive updates and mergeable updates, resp. Let  $B = (B_1 \dots B_n)$  be a model tuple in the carrier, then we define

$$\text{Corr}(\ell)(B) = \{(R, R_1 \dots R_n) \mid \exists A = (A_1 \dots A_n) \in \prod \text{Ob}(\mathbf{A}) : R \in \text{Corr}(\kappa)(A), R_i \in \text{Corr}(b_i)(B_i, A_i)\}$$

[Consistent] corrs are exactly the tuples  $(R, R_1 \dots R_n)$ , in which  $R$  and all  $R_i$  are [consistent] corrs in their respective multispaces. This yields  $\partial_i^\ell(R, R_1 \dots R_n) = \partial_{\mathbf{B}}^{b_i} R_i$  (see Fig.10). Propagation operations are defined as compositions of consecutive lens' executions as described below (we will use the dot notation for operation application and write  $x.\text{op}$  for  $\text{op}(x)$ , where  $x$  is an argument).

Given a model tuple  $(B_1 \dots B_n) \in \mathbf{B}_1 \times \dots \times \mathbf{B}_n$ , a corr  $(R, R_1 \dots R_n)$ , and update  $v_i: B_i \rightarrow B'_i$  in  $\text{Ar}(\mathbf{B}_i)$ , we define, first for  $j \neq i$ ,

$$v_i.\ell.\text{ppg}_{ij}^{(R, R_1 \dots R_n)} \stackrel{\text{def}}{=} v_i.(b_i.\text{ppg}_{\mathbf{B}\mathbf{A}}^{R_i}).(\kappa.\text{ppg}_{ij}^R).(b_j.\text{ppg}_{\mathbf{A}\mathbf{B}}^{R_j}),$$

and then  $v_i.\ell.\text{ppg}_{ii}^{(R, R_1 \dots R_n)} \stackrel{\text{def}}{=} v_i.b_i.\text{ppg}_{\mathbf{B}\mathbf{B}}^{R_i}$ . Note that all internal amendments to  $u_i = v_i.(b_i.\text{ppg}_{\mathbf{B}\mathbf{A}}^{R_i})$  produced by  $\kappa$ , and to  $u'_j = u_i.(\kappa.\text{ppg}_{ij}^R)$  produced by  $b_j$ , are identities due to the Junction conditions. This allows us to set corrs properly and finish propagation with the three steps above:  $v_i.\ell.\text{ppg}_{i\star}^{(R, R_1 \dots R_n)} \stackrel{\text{def}}{=} (R', R'_1 \dots R'_n)$  where  $R' = u_i.\kappa.\text{ppg}_{i\star}^R$ ,  $R'_j = u'_j.b_j.\text{ppg}_{\mathbf{A}\star}^{R_j}$  for  $j \neq i$ , and  $R'_i = v_i.b_i.\text{ppg}_{\mathbf{B}\star}^{R_i}$ . We thus have a lens  $\ell$  denoted by  $\kappa^*(b_1, \dots, b_n)$ .  $\square$

**Theorem 3 (Star Composition).** *Given a star configuration of lenses as above. Let all underlying model spaces be well-behaved, cf. Def. 2, and the Junction conditions hold. Then the following holds: If lens  $\kappa$  and all lenses  $b_i$  are (very) wb, then  $\kappa^*(b_1, \dots, b_n)$  is also (very) wb.*

*Proof.* Fulfilment of **Stability** is obvious. **Reflect1-3** follow immediately from the definition of  $\ell.\text{ppg}_{ii}$  above, since the first step of the above propagation procedure already enjoy sequential compatibility and idempotency by **Reflect1-3** for  $b_i$ . This proves the wb part of the theorem.

Now we prove that the composed lens is very wb if the components are such. If all model spaces are well-behaved, then operation **merge** preserves identities, cf. Def. 2. Then **KPutPut** for any lens  $\kappa$  reduces to the simplified form  $\text{ppg}_{ij}^R(u_i; v_i) = \text{ppg}_{ij}(u_i); \text{ppg}_{ij}(v_i)$ , if  $u_i$  and  $v_i$  are  $\kappa$ -closed. The Junction condition and the fact that operation **ppg** preserves the **K $\blacktriangleright$** -property, guarantee that propagation of composed updates from  $\mathbf{A}_i$  to  $\mathbf{A}_j$  and further to  $\mathbf{B}_j$  is not disturbed by reflective updates. Hence, **KPutPut** carries over from  $b_i$  to  $\kappa^*(b_1, \dots, b_n)$ .  $\square$

### 6.3 Lens composition and Invertibility

Unfortunately, even if all component lenses are invertible, the composed star-lens is not necessarily such as we will show in the next subsection. In subsection 6.3.2, we discuss a seemingly counter-example to this negative result and show that the state-based setting for update propagation can be confusing.

**6.3.1 Counter-example.** Consider a class of simple model spaces, whose objects (called models) are natural numbers plus some fixed symbol  $\perp$  denoting an undefined number. Thus, for a space  $\mathbf{A}$  from this class, we assume  $\text{Ob}(\mathbf{A}) = \{\perp\} \cup \mathbf{A}!$  with  $\mathbf{A}! = \{A_1, A_2, \dots\}$  being a set of natural numbers:  $A_i = \{i\}$ . Models form the set  $\mathbf{A}!$  are called *certain* while models  $\perp$  are *uncertain*. Updates are all possible pairs of models:  $\text{Ar}(\mathbf{A}) = \text{Ob}(\mathbf{A}) \times \text{Ob}(\mathbf{A})$  (such categories are often called co-discrete or chaotic), that is,

$$\text{Ar}(\mathbf{A}) = \{(\perp, \perp)\} \cup \{\perp\} \times \mathbf{A}! \cup \mathbf{A}! \times \{\perp\} \cup \mathbf{A}! \times \mathbf{A}!$$

Now consider three model spaces of the type specified above: space  $\mathbf{B}_1$  with certain models  $B_1 = \{0, 2, 6\}$ , space  $\mathbf{A}$  with certain models  $A = \{2, 5, 6\}$  and space  $\mathbf{B}_2$  with  $B_2 = \{4, 7\}$ . For all these spaces, mergeable pairs are only those with at least one identity and the same necessarily holds for  $\mathbf{K}^{\blacktriangleright}$ , hence all spaces are well-behaved (cf. Def. 2).

Suppose we have a lens  $b_1$  over spaces  $\mathbf{B}_1$  and  $\mathbf{A}$  (see Fig. 11). The set of corrs for a pair  $(b, a) \in \text{Ob}(\mathbf{B}_1) \times \text{Ob}(\mathbf{A})$  is the singleton set  $\{(b, a)\}$ . A corr is consistent, if and only if  $b, a$  are both certain or both uncertain, i.e., consistency amounts to equal certainty. Then to restore consistency, updates of type  $B_1 \times \{\perp\}$  are to be mapped to similar updates of type  $A \times \{\perp\}$ , and updates of type  $B_1 \times B_1$  can be mapped to the corresponding identity updates on the  $A \times A$  as they do not destroy consistency. The propagation operation from  $\mathbf{A}$  to  $\mathbf{B}_1$  of the types above are defined similarly.

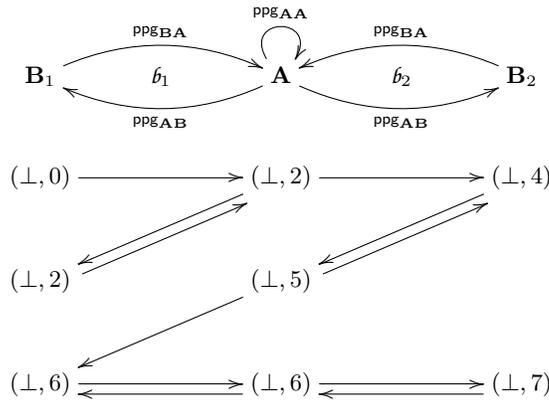


Fig. 11: Schematic description of the counterexample

Now we need to define how to propagate updates of the type  $\{\perp\} \times B_1$  over the corr  $\{(\perp, \perp)\}$ .

This is specified in Fig. 11 by arrows going from the left-most column of updates to the middle column. The idea is to find in  $\mathbf{A}$  the nearest model so that updating  $\perp$  to 0 or 2 in  $\mathbf{B}_1$  goes to updating  $\perp$  to 2 in  $\mathbf{A}$ , while updating  $\perp$  to 6 goes to updating  $\perp$  to 6. In the similar way, propagation of updates from  $\{\perp\} \times A$  to  $\{\perp\} \times B_1$  is defined as shown by arrows from the middle to the left column. This defines lens  $\mathfrak{b}_1$ , and it is easy to see it is well-behaved and, moreover, invertible. The latter is established by the direct examination of all possible  $\text{ppg}$ -compositions, e.g.,

$$\begin{aligned} (0, \perp) \cdot \text{ppg}_{\mathbf{BA}}^{(0,2)} \cdot \text{ppg}_{\mathbf{AB}}^{(0,2)} &= (0, \perp) \text{ (even strong invertibility holds)} \\ (0, 2) \cdot \text{ppg}_{\mathbf{BA}}^{(0,2)} \cdot \text{ppg}_{\mathbf{AB}}^{(0,2)} &= (0, 0) \neq (0, 2) \text{ but } (0, 2) \cdot \text{ppg}_{\mathbf{BA}}^{(0,2)} = (0, 0) \cdot \text{ppg}_{\mathbf{BA}}^{(0,2)} \\ (\perp, 0) \cdot \text{ppg}_{\mathbf{BA}}^{(\perp, \perp)} \cdot \text{ppg}_{\mathbf{AB}}^{(\perp, \perp)} &= (\perp, 2) \neq (\perp, 0) \text{ but } (\perp, 0) \cdot \text{ppg}_{\mathbf{BA}}^{(\perp, \perp)} = (\perp, 2) \cdot \text{ppg}_{\mathbf{BA}}^{(\perp, \perp)} \end{aligned}$$

(the last example is specific for the lens  $\mathfrak{b}_1$  specified in Fig. 11). Similarly, we define a wb lens  $\mathfrak{b}_2$  as shown in Fig. 11 and check it is also invertible.

Space  $\mathbf{A}$  in the middle can be extended to a trivially wb and invertible identity lens  $\mathfrak{k} = \mathfrak{id}_n(\mathbf{A})$  with  $n = 2$  (cf. Example 3).

Now, as the Junction condition trivially holds for the triple  $(\mathfrak{b}_1, \mathfrak{k}, \mathfrak{b}_2)$ , we obtain a star lens  $\ell := \mathfrak{k}^*(\mathfrak{b}_1, \mathfrak{b}_2)$  composed from wb invertible lenses. However, this lens is not invertible as the following computation shows:

$$(\perp, 0) \cdot (\ell \cdot \text{ppg}_{12}) \cdot (\ell \cdot \text{ppg}_{21}) \cdot (\ell \cdot \text{ppg}_{12}) = (\perp, 7) \neq (\perp, 4) = (\perp, 0) \cdot (\ell \cdot \text{ppg}_{12})$$

where we omitted the upper scripts  $(\perp, \perp)$  near  $\text{ppg}$ -symbols (recall that  $\ell \cdot \text{ppg}_{12}$  is defined as the sequential composition  $(\mathfrak{b}_1 \cdot \text{ppg}_{\mathbf{BA}}) \cdot (\mathfrak{k} \cdot \text{ppg}_{\mathbf{AA}}) \cdot (\mathfrak{b}_2 \cdot \text{ppg}_{\mathbf{AB}})$  and similarly for  $\ell \cdot \text{ppg}_{12}$ ).

**Theorem 4.** *Star-composition does not preserve invertibility*

**6.3.2 Invertibility and (binary) state-based symmetric lenses with complement (ssc-lenses): a long standing confusion.** The example above may seem to be contradicting to paper [18], where symmetric lenses are studied in the state-based setting as ssc-lenses. In that paper, an invertibility law called *round-tripping* is required for any ssc-lens, and it is proved that sequential composition of such lenses is again an ssc-lens and hence enjoys round-tripping. In paper [11], we show that an ssc-lens is nothing but a symmetric delta lens over co-discrete model spaces (see also [21]), i.e., exactly a lens of the type we have considered in the counter-example above. Moreover, in the star-composition instance we have considered, lens  $\mathfrak{k} = \mathfrak{id}_2(\mathbf{A})$  plays a dummy role and, in fact, we have dealt with sequential composition of two ssc-lenses  $\mathfrak{b}_1; \mathfrak{b}_2$  as defined in [18]. Then, how could it happen that our composed lens does not satisfy (even weak) invertibility, while the corresponding result in [18] asserts that the composed lens must satisfy the seemingly stronger roundtripping law?

The source of confusion is the state-based setting for update propagation, in which a law that looks like demanding strong invertibility is actually a simple Stability law demanding identity preservation. Indeed, in [18, Def.2.1 on p.2], they define a binary symmetric lens over state spaces  $X$  and  $Y$  with propagation operations

$$(3) \quad \text{ppg}_{XY}: X \times \text{Corr} \rightarrow Y \times \text{Corr} \text{ and } \text{ppg}_{YX}: X \times \text{Corr} \leftarrow Y \times \text{Corr}$$

(they call elements of set  $\mathbf{Corr}$  complements rather than corrs, but as it is shown in [11,21], the two notions are equivalent). The law (PutRL) called *round-tripping* is defined thusly: for any states  $x$  and  $y$  (in our notation,  $A'_1$  and  $A'_2$ ) and complements  $c, c'$  (i.e., our corrs,  $R, R'$ ), the following condition holds ([18, Def.2.1 on p.2]):

$$(4) \quad \mathbf{ppg}_{XY}(x, c) = (y, c') \text{ implies } \mathbf{ppg}_{YX}(y, c') = (x, c)$$

Figure 12 specifies the story described by the two equations in (4) diagrammatically: the upper square specifies data of the left equation, and the lower square specifies the right equation. Now it is seen that the right equation specifies the Stability law of delta lenses, while the left equation is used to ensure that states  $x$  and  $y$  are related by complement/corr  $c'$ . The left equation is needed as the state-based lens framework does not have primitives to say that states  $(x, y)$  are consistent via corr  $c'$  (which is done in delta lenses by boundary functions  $\partial_i$ ), and thus they encode this fact equationally, which together with the right equation creates a false impression of a round-tripping law. Indeed, equations (4) have nothing to do with round-tripping: to describe the latter, we should specify how the result of  $\mathbf{ppg}_{YX}(y, c)$  is related to  $x$  rather than how  $\mathbf{ppg}_{YX}(y, c')$  is related to  $x$ .

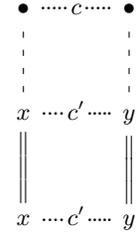


Fig. 12: “Round-tripping” story

The main problem is not in that Stability is called Round-tripping, the problem is that the issue of an actual round-tripping law and its preservation under lens composition is not stated in the state-based lens framework because it is (fallaciously) considered solved! The counter-example presented above makes the problem even more challenging as even weak round-tripping is not preserved under lens composition.

### 6.4 Sequential Lego 2: Assembling $n$ -ary Lenses from Binary Lenses

This section shows how to assemble  $n$ -ary (symmetric) lenses from binary asymmetric lenses modelling view computation [10]. As the latter is a typical  $\mathbf{bx}$ , the well-behavedness of asymmetric lenses has important distinctions from well-behavedness of general (symmetric  $\mathbf{mx}$ -tailored) lenses.

**Definition 13 (Asymmetric Lens, cf. [10]).** An *asymmetric lens (a-lens)* is a tuple  $\mathcal{L} = (\mathbf{A}, \mathbf{B}, \mathbf{get}, \mathbf{put})$  with  $\mathbf{A}$  a model space called the (*abstract*) *view*,  $\mathbf{B}$  a model space called the *base*,  $\mathbf{get} : \mathbf{A} \leftarrow \mathbf{B}$  a functor (read “get the view”), and  $\mathbf{put}$  a family of operations  $(\mathbf{put}^B \mid B \in \mathbf{Ob}(\mathbf{B}))$  (read “put the view update back”) of the following arity. Provided with a view update  $v : \mathbf{get}(B) \rightarrow A'$  at the input, operation  $\mathbf{put}^B$  outputs a base update  $\mathbf{put}_B^B(v) = u' : B \rightarrow B''$  and an amendment to the view update,  $\mathbf{put}_B^B(v)$  or  $v^\circledast : A' \rightarrow A''$ .

A view update  $v : \mathbf{get}(B) \rightarrow A'$  is called *closed* if  $v^\circledast = \text{id}_{A'}$ , and an a-lens is *closed* if all its view updates are closed. □

The following is a specialization of Def. 7, in which consistency between a base  $B$  and a view  $A$  is understood as equality  $\mathbf{get}(B) = A$ .

**Definition 14 (Well-behavedness).** An a-lens is *well-behaved (wb)* if it satisfies the following laws for all  $B \in \mathbf{Ob}(\mathbf{B})$  and  $v : \mathbf{get}(B) \rightarrow A'$

- (Stability) if  $v = \text{id}_{\text{get}(B)}$ , then  $\text{put}_b^B(v) = \text{id}_B$  and  $v^\circledast = \text{id}_{\text{get}(B)}$   
 (Reflect0) if  $v = \text{get}(u)$  for some  $u: B \rightarrow B'$ , then  $v^\circledast = \text{id}_{A'}$   
 (Reflect1)  $(v, v^\circledast) \in \mathbf{K}_A^\blacktriangleright$   
 (Reflect2)  $\text{put}_b^B(v; v^\circledast) = \text{put}_b^B(v)$   
 (PutGet)  $v.\text{put}_b^B.\text{get} = v; v^\circledast$  (the dot notation is used to highlight the name of the law).  $\square$

In contrast to the general lens case, a wb a-lens features Reflect0 — a sort of a self-Hippocratic law regulating the necessity of amendments: the latter is not allowed as soon as consistency can be restored without amendment. Another distinction is inclusion of an invertibility law PutGet into the definition of well-behavedness: PutGet together with Reflect2 provide (weak) invertibility:  $\text{put}_b^B(\text{get}(\text{put}_b^B(v))) = \text{put}_b^B(v)$ . Reflect3 is omitted as it is implied by Reflect0 and PutGet.

Any a-lens  $\mathfrak{b}^\leftarrow = (\mathbf{A}, \mathbf{B}, \text{get}, \text{put})$  gives rise to a binary symmetric lens  $\mathfrak{b}$ . Its carrier consists of model spaces  $\mathbf{A}$  and  $\mathbf{B}$ , for which we assume  $\mathbf{K}^\blacktriangleright$  and  $\mathbf{K}^\blacktriangleleft$  to be defined such that the spaces are well-behaved, cf. Def. 2. The set of corrs of a pair  $(A, B) \in \text{Ob}(\mathbf{A}) \times \text{Ob}(\mathbf{B})$  is the singleton  $\{(A, B)\}$  with  $\partial_{\mathbf{A}}(A, B) = A$  and  $\partial_{\mathbf{B}}(A, B) = B$ , and consistent corrs are exactly those  $(A, B)$ , for which  $A = \text{get}(B)$ .

For a consistent corr  $\widehat{B} = (\text{get}(B), B)$ , we need to define six operations  $\mathfrak{b}.\text{ppg}_{\_ \_}^{\widehat{B}}$ . Below we will write the upper index as  $B$  rather than  $\widehat{B}$  to ease the notation. For a view update  $v: A \rightarrow A'$  and the corr  $\widehat{B}$ , we define

$$\text{ppg}_{\mathbf{A}\mathbf{B}}^B(v) = \text{put}_b^B(v): B \rightarrow B'', \text{ppg}_{\mathbf{A}\mathbf{A}}^B(v) = \text{put}_v^B(v): A' \rightarrow A'', \text{ppg}_{\mathbf{A}\star}^B(v) = \widehat{B}''$$

The condition  $A'' = \text{get}(B'')$  for  $\mathfrak{b}^\leftarrow$  means that  $\widehat{B}''$  is again a consistent corr with the desired boundaries. For a base update  $u: B \rightarrow B'$  and the corr  $\widehat{B}$ , we define

$$\text{ppg}_{\mathbf{B}\mathbf{A}}^B(u) = \text{get}(u), \text{ppg}_{\mathbf{B}\mathbf{B}}^B(u) = \text{id}_{B'}, \text{ppg}_{\mathbf{B}\star}^B(u) = \widehat{B}'$$

Functoriality of  $\text{get}$  yields consistency of  $\widehat{B}'$ .

**Definition 15 (Very well-behaved a-lenses).** A wb a-lens is called *very well behaved (very wb)*, if the corresponding binary symmetric lens is such. cf. Def. 9. Of course, this definition restricts the behaviour only for the  $\text{put}$ -part of the a-lens (hence the name of the law), since functor  $\text{get}$  is compositional anyway.  $\square$

**Lemma 5.** *Let  $\mathfrak{b}^\leftarrow$  be a (very) wb a-lens and  $\mathfrak{b}$  the corresponding symmetric lens. Then all base updates of  $\mathfrak{b}$  are closed, and  $\mathfrak{b}$  is (very) wb and invertible.*

*Proof.* Base updates are closed by the definition of  $\text{ppg}_{\mathbf{B}\mathbf{B}}^B$ . Well-behavedness follows from wb-ness of  $\mathfrak{b}^\leftarrow$ . Invertibility has to be proved in two directions:  $\text{ppg}_{\mathbf{B}\mathbf{A}}^B; \text{ppg}_{\mathbf{A}\mathbf{B}}^B; \text{ppg}_{\mathbf{B}\mathbf{A}}^B = \text{ppg}_{\mathbf{B}\mathbf{A}}^B$  follows from (PutGet) and (Reflect0), the other direction follows from (PutGet) and (Reflect2), see the remark after Def.14. The “vwb”-implication follows directly from Def.15.  $\square$

*Remark 2 (A-lenses as view transformation engines).* In terms of Varró *et al* [28], a wb a-lens can be seen as an abstract algebraic model of a bidirectional view transformation engine. This engine is assumed to be a) consistent: it is

a basic lens law that propagation must always produce a consistent corr, b) incremental, as lenses propagate changes rather than create the target models from scratch, c) validating, as the result of propagation is always assumed to satisfy the target metamodel constraints (either by considering the entire model space to be only populated by valid models, or by introducing a subclass of valid models and require the result of the propagation to get into that subclass). Reactiveness, i.e., whether the engine executes on-demand or in response to changes, is beyond the lens framework. To address this and similar concerns, we need a richer framework of lenses augmented with an *organizational* structure introduced in [6] (see also [8] for a concise presentation).

**Theorem 5 (Lenses from Spans).** *An  $n$ -ary span of (very) wb asymmetric lenses  $b_i^{\leftarrow} = (\mathbf{A}_i, \mathbf{B}, \text{get}_i, \text{put}_i)$ ,  $i = 1..n$  with a common base  $\mathbf{B}$  of all  $b_i^{\leftarrow}$  gives rise to a (very) wb symmetric lens denoted by  $\Sigma_{i=1..n}^{\mathbf{B}} b_i^{\leftarrow}$ .*

*Proof.* An  $n$ -ary span of a-lenses  $b_i^{\leftarrow}$  (all of them interpreted as symmetric lenses  $b_i$  as explained above) is a construct equivalent to the star-composition of Def. 6.2.3, in which lens  $\kappa = \text{id}_n(\mathbf{B})$  (cf. Ex. 3) and peripheral lenses are lenses  $b_i$ . The junction condition is satisfied as all base updates are  $b_i$ -closed for all  $i$  by Lemma 5, and also trivially closed for any identity lens. The theorem thus follows from Theorem 3. The “very wb”-part follows directly from Def. 15. Note that a corr in  $\text{Corr}(\Sigma_{i=1..n}^{\mathbf{B}} b_i^{\leftarrow})$  is nothing but a single model  $B \in \text{Ob}(\mathbf{B})$  with boundaries being the respective  $\text{get}_i$ -images.  $\square$

The theorem shows that combining a-lenses in this way yields an  $n$ -ary symmetric lens, whose properties can automatically be inferred from the binary a-lenses.

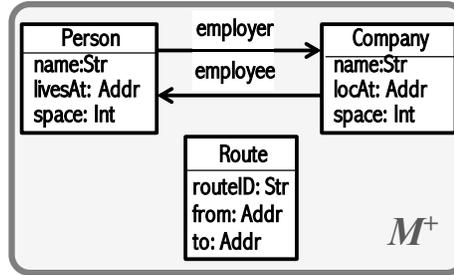


Fig. 13: Merged Metamodel

*Running example.* Figure 13 shows a metamodel  $M^+$  obtained by merging the three metamodels  $M_{1,2,3}$  from Fig. 2 without loss and duplication of information. In addition, for persons and companies, the identifiers of their original model spaces can be traced back via attribute “space” (Route-objects are known to appear in space  $\mathbf{A}_3$  and hence do not need such an attribute). As shown in [26], any consistent multimodel  $(A_1 \dots A_n, R)$  can be merged into a comprehensive model  $A^+$  instantiating  $M^+$ . Let  $\mathbf{B}$  be the space of such together with their comprehensive updates  $u^+ : A^+ \rightarrow A'^+$ .

For a given  $i \leq 3$ , we can define the following a-lens  $b_i^{\leftarrow} = (\mathbf{A}_i, \mathbf{B}, \text{get}_i, \text{put}_i)$ :  $\text{get}_i$  takes update  $u^+$  as above and outputs its restriction to the model containing only objects recorded in space  $\mathbf{A}_i$ . Operation  $\text{put}_i$  takes an update  $v_i : A_i \rightarrow A'_i$  and first propagates it to all directions as discussed in Sect. 3, then merges these

propagated local updates into a comprehensive  $\mathbf{B}$ -update between comprehensive models. This yields a span of a-lenses that implements the same synchronization behaviour as the symmetric lens discussed in Sect. 3.

**From (symmetric) lenses to spans of a-lenses.** There is also a backward transformation of (symmetric) lenses to spans of a-lenses. Let  $\ell = (\mathcal{A}, \text{ppg})$  be a (very) wb lens. It gives rise to the following span of (very) wb a-lenses  $\ell_i^{\leftarrow} = (\partial_i(\mathcal{A}), \mathbf{B}, \text{get}_i, \text{put}_i)$  where space  $\mathbf{B}$  has objects  $(A_1 \dots A_n, R)$  with  $R \in \text{Corr}(A_1 \dots A_n)$  a consistent corr, and arrows are the update tuples of  $\mathcal{A}$ . Functors  $\text{get}_i : \mathbf{B} \rightarrow \mathbf{A}_i$  are projection functors. Given  $B = (A_1 \dots A_n, R)$  and local update  $u_i : A_i \rightarrow A'_i$ , let

$$\text{put}_{ib}^B(u_i) \stackrel{\text{def}}{=} (u'_1, \dots, u'_{i-1}, (u_i; u'_i), u'_{i+1}, \dots, u'_n) : (A_1 \dots A_n, R) \rightarrow (A'_1 \dots A'_n, R'')$$

where  $u'_j \stackrel{\text{def}}{=} \text{ppg}_{ij}^R(u_i)$  (all  $j$ ) and  $R'' = \text{ppg}_{i\star}^R(u_i)$ . Finally,  $\text{put}_{iv}^B(v_i) \stackrel{\text{def}}{=} \text{ppg}_{iv}^R(v_i)$ , i.e.,  $v_i^{\textcircled{B}} = v_i^{\textcircled{R}}$ . Validity of Stability, Reflect0-2, PutGet and KPutput law directly follows from the above definitions. This yields

**Theorem 6.** *Let  $\ell = (\mathcal{A}, \text{ppg})$  be a (very) wb symmetric lens. Then the multi-span*

$$(\ell_i^{\leftarrow} = (\partial_i(\mathcal{A}), \mathbf{B}, \text{get}_i, \text{put}_i) \mid i = 1..n)$$

*of (very) wb a-lenses (where  $\mathbf{B}$  is defined as specified above) has the same synchronisation behaviour w.r.t. model spaces  $(\partial_i(\mathcal{A}) \mid i = 1..n)$ .  $\square$*

An open question is whether the span-to-lens transformation in Thm.5 and the lens-to-span transformation of Thm. 6 are mutually inverse. It is easy to see that the chain

$$(5) \quad \text{symmetric lens} \rightarrow \text{wide span} \rightarrow \text{symmetric lens}$$

results in a lens isomorphic to the initial one, but the chain

$$(6) \quad \text{wide span} \rightarrow \text{symmetric lens} \rightarrow \text{wide span}$$

ends with a different span as the first transformation above loses information about updates in the head of the span. As shown by Johnson and Rosebrugh in [22] for the binary case, the two spans can only be equivalent modulo certain equivalence relation, and an equivalence relation between lenses is also needed to align all constructs together. These equivalences may be different for our multiary lenses with amendments, and we leave this important question for future research.

## 7 Future work

We list and briefly comment on several important and challenging tasks for the multiary lens framework.

### 7.1 Categorification of corrs

A distinctive feature of the framework developed in the paper is the triviality of the corr updates — they are just pairs  $(R, R')$  of the old and the new corr (see Remark 1 on p.15). However, in practice, new corrs would be computed incrementally with deltas rather than afresh. To make the framework closer to

practice, we need to change the notion of a multimodel update  $\mathbf{u}: \mathcal{A} \rightarrow \mathcal{A}'$  and consider it to be a pair  $\mathbf{u} = (u, r)$  with  $u = (u_1..u_n)$  a feet update and  $r: R \rightarrow R'$  a corr update rather than just a pair of states  $(R, R')$ . Then we would obtain a setting based on a category  $\mathcal{R}$  of multimodels (that includes both local updates and corr updates) together with boundary projection functors  $\partial_i: \mathcal{R} \rightarrow \mathbf{A}_i$ ,  $i = 1..n$ , which take a multimodel update  $\mathbf{u} = (u_1..u_n, r)$  and select its corresponding component,  $\partial_i(\mathbf{u}) = u_i$ .<sup>9</sup>

For multiary update propagation, we require each projection  $\partial_i$  to be the get-part of an asymmetric lens (with amendment) that for a given multimodel state  $R$ , puts any foot update  $u_i: A_i \rightarrow A'_i$ ,  $A_i = \partial_i(R)$  back to a multimodel update  $\mathbf{u} = \mathbf{put}_i^R(u_i): \mathcal{A} \rightarrow \mathcal{A}'$ . Note that as  $\mathbf{u}$  includes all feet updates, operation  $\mathbf{put}_i^R$  actually provides all local propagation operations  $\mathbf{ppg}_{ij}$  we considered in the paper:  $\mathbf{ppg}_{ij}^R(u_i) = \partial_j(\mathbf{put}_i^R(u_i))$ . The Putget law is the equality  $\partial_i(\mathbf{u}) = u_i; u_i^\circledast$ . In this way, a multiary symmetric lens is—by definition—a multiary span of asymmetric lenses  $(\partial_i, \mathbf{put}_i): \mathcal{R} \rightarrow \mathbf{A}_i$ ,  $i = 1..n$ .

Besides being better aligned with practice, the setting above would simplify notation and probably some technicalities, but its advantages were recognized when the paper was already written and submitted for reviewing. With a great regret, an accurate theory of the categorified version of multiary lenses is thus left for future work.

**7.2 Instantiation of the framework.** Although we presented several simple examples of multiary wb lenses, having practically interesting examples, or even better, a pattern for generating practically interesting examples, is an extremely important task. We think that the GDG framework for synchronization [33] is a promising foundation for such work, and it has actually already began in [34] (see the next section for more detailed comments).

**7.3 Concurrent updates.** As discussed in section Sect. 2.4, extending the lens formalism to accommodate concurrent updates is both useful and challenging. One of the main problems to solve is how to reconcile non-determinism inherent in conflict resolution strategies and serialization of update propagation with the deterministic nature of the lens framework.

**7.4 Richer compositional framework.** There are several open issues here even for the non-concurrent case (not to mention its future concurrent generalization). First, our pool of lens composition constructs is far incomplete: we need to enrich it, at least, with (i) sequential composition of a-lenses with amendment so that a category of a-lenses could be built, and (ii) a relational composition of symmetric lenses sharing several of their feet (similar to relational join). It is also important to investigate composition with weaker junction conditions than that we considered.

<sup>9</sup> Given some construct  $X$  depending on set  $S$ , the passage to a setting in which  $S$  is a non-trivial category is often called *categorification* of  $X$ , hence, the title of the subsection. The quality of being a *non-trivial* category is essential. A set  $S$  can be seen as a category  $\mathbf{S}$  in two ways: *discrete*, when the only  $\mathbf{S}$  arrows are identities, and *co-discrete* or *chaotic*, in which for any pair of  $S$ 's elements  $s, s'$ , there is one and only one arrow  $(s, s'): s \rightarrow s'$ . If  $s' = s$ , this arrow is the identity of  $s$ . We refer to both such categories  $\mathbf{S}$  as trivial.

**7.5 Invertibility.** Invertibility nicely fits in some but not all of our results. It is a sign that we do not well understand the nature of invertibility. Perhaps, while invertibility is essential for  $\text{bx}$ , its role for  $\text{mx}$  may be less important.

**7.6 Hippocraticness.** A very natural Hippocraticness requirement (introduced by Stevens for  $\text{bx}$  in [31]) may have less weight in the  $\text{mx}$  world. Indeed, the examples we considered in the paper show that non-Hippocratic consistency restoration may be practically reasonable and useful. A further research in this direction is needed.

## 8 Related Work

Compositionality as a fundamental principle for building synchronization tools was proposed by Pierce and his coauthors in [15,1] for the state-based asymmetric case, and further developed for the binary symmetric case by Hofmann *et al.* in [18], for the state-based setting, and in [19], for a specific delta-based setting (where deltas are understood operationally as edits). The deficiencies of the state-based framework are discussed in detail in [10,11] (see also our Sect. 6.3.2), where the notion of asymmetric and symmetric binary delta lenses were proposed, but lens composition was not specially considered (except easy sequential composition of  $\text{a-lenses}$  in [10]). In section 6.1 of [4], several results for delta lens composition are considered, including a special case of the star composition, when two asymmetric lenses form a cospan, but the consistency relation at the apex of this cospan is not assumed to be an equality and, hence, to be maintained by a general binary symmetric lens. A fundamental theory of composing (and decomposing) binary symmetric lenses from (into) spans or cospans of asymmetric ones was developed by Johnson and Rosebrugh [22,23]; the theory is based on equivalences of lenses w.r.t. their behaviour. In all these works, neither amendments nor intelligent (and hence constrained) versions of Putput are considered (and only binary lenses are taken into consideration).

In the turn from the binary to the multiary update propagation, the work by Trollmann and Albayrak [33,34,35] is the closest to ours conceptually and in part technically. We made several brief remarks in the introduction, and now can provide more details. Trollmann and Albayrak begin with defining a class  $\mathcal{S}$  of diagram shapes (they say, bases), and a multimodel is a graph diagram of a certain shape  $S \in \mathcal{S}$  in a adhesive category  $\mathbf{G}$  intended to model attributed typed graphs and their morphisms, i.e., a multimodel is a functor  $M: S \rightarrow \mathbf{G}$  with  $S \in \mathcal{S}$ . A multimodel update  $m: M \rightarrow M'$  is a span of natural transformations,  $m_d: M \leftarrow \hat{m}$ ,  $m_a: \hat{m} \rightarrow M'$ , whose head  $\hat{m}$  is a diagram of the same shape, and injective transformations  $m_d$  and  $m_a$  specify, resp., deletions and additions provided by update  $m$ . For a fixed  $S$ , this gives us a non-trivial category of multimodels  $\mathbf{M}(S)$  to be compared with the categorified version of our framework based on category  $\mathbf{Corr}$  with arrows as described in Sect. 7.1. A major distinction between the two is that  $\mathbf{Corr}$  is a general model space category without any further restrictions, everything needed is provided by a family of boundary functors  $\partial_i: \mathbf{Corr} \rightarrow \mathbf{A}_i$ . In contrast, although the class of categories  $\{\mathbf{M}(S) \mid S \in \mathcal{S}\}$  is broad enough to be practically interesting, it does not include some corr structures appearing in applications, e.g., such as in our Running example, or in paper [9] focused on UML modelling, or in our paper [26], in which a very general pattern for corrs is described as a partial span of

graph morphisms, or, finally, in general rule-based synchronization engines as described in [3]. An advantage of the abstract lens framework is that all these constructs are uniformly modelled as category **Corr**.

Further distinctions appear when we consider how update propagation operations are defined. Trollmann and Albayrak consider the concurrent case (a truly impressive achievement), which we did not approach yet. However, they consider only two basic laws (Stability and Correctness) borrowed from the binary case, while our repertoire of laws is richer. They also have a sort of amendment operation implemented by Consistency Creating Rules, but it is a purely local operation independent on any other model in the multimodel. In contrast, our amendments are tightly related to other models and in this sense are (self) propagation operations. Finally, a major distinction is that in their setting, consistency restoration is non-deterministic while our lenses are deterministic. Overall, lenses appear as an abstract algebraic interface to update propagation that can be implemented by different ways, e.g., with GDG or by repair rules as in [3,30].

For the state-based lens setting, the work closest in spirit to the turn from binary to multiary lenses, is Stevens’ paper [32]. Her and our goals are similar, but the technical realizations are different even besides the state- vs. delta-based opposition. Stevens works with restorers, which take a multimodel (in the state-based setting, just a tuple of models) presumably *inconsistent*, and restores consistency by changing some models in the tuple while keeping other models (from the *authority set*) unchanged. In contrast, lenses take a *consistent* multimodel *and* updates, and return a consistent multimodel and updates. As we argued in Sect. 3, including updates into the input data of the restoration operation allows better “tuning” update propagation policies as the inherited uncertainty of consistency restoration is reduced. Another important difference is update amendments, which are not considered in [32] – models in the authority set are intact. Yet another distinction is how the multiary vs. binary issue is treated. Stevens provides several results for decomposing an  $n$ -ary relation  $R \in \text{Corr}(\mathbf{A})$  into binary relations  $R_{ij} \subseteq \text{Ob}(\mathbf{A}_i) \times \text{Ob}(\mathbf{A}_j)$  between the components. For us, a relation is inherently  $n$ -ary, i.e., a set  $R$  of  $n$ -ary links endowed with an  $n$ -tuple of projections  $\partial_i: R \rightarrow A_i$  uniquely identifying links’ boundaries. Thus, while Stevens considers “binarization” of a relation  $R$  by a *chain* of binary relations over the “perimeter”  $A_1 \dots A_n$ , we binarize it via the corresponding *span* of (binary) mappings  $(\partial_1, \dots, \partial_n)$  (UML could call this process *reification*). Our (de)composition results demonstrate advantages of the span view.

Discussion of several other works in the state-based world, notably by Macedo *et al* [27] can be found in [32]. Several remarks about the related work on the Putput law have already been made in Sect. 2.3

## 9 Conclusion

Multimodel synchronization is an important practical problem, which cannot be fully automated but even partial automation would be beneficial. A major problem in building an automatic support is uncertainty inherent to consistency restoration. In this regard, restoration via update propagation rather than immediate repairing of an inconsistent state of the multimodel has an essential advantage: having the update causing inconsistency as an input for the restoration operation can guide the propagation policy and essentially reduce the uncertainty. We thus come to the scenario of multiple model synchronization via

multi-directional update propagation. We have also argued that reflective propagation to the model whose change originated inconsistency is a reasonable feature of the scenario.

We presented a mathematical framework for synchronization scenarios as above based on a multiary generalization of binary symmetric delta lenses introduced earlier, and enriched it with reflective propagation and KPutput law ensuring compatibility of update propagation with update composition in a practically reasonable way (in contrast to the strong but unrealistic Putput). We have also defined several operations of composing multiary lenses in parallel and sequentially. Our lens composition results make the framework interesting for practical applications: if a tool builder has implemented a library of elementary synchronization modules based on lenses and, hence, ensuring basic laws for change propagation, then a complex module assembled from elementary lenses will automatically be a lens and thus also enjoys the basic laws. This allows one to avoid additional integration testing, which can essentially reduce the cost of synchronization software.

**Acknowledgement.** We are really grateful to anonymous reviewers for careful reading of the manuscript and detailed, pointed, and stimulating reviews, which essentially improved the presentation and discussion of the framework.

## References

1. Bohannon, A., Foster, J.N., Pierce, B.C., Pilkiewicz, A., Schmitt, A.: Boomerang: resourceful lenses for string data. In: Necula, G.C., Wadler, P. (eds.) POPL. pp. 407–419. ACM (2008)
2. Chechik, M., Nejati, S., Sabetzadeh, M.: A relationship-based approach to model integration. *ISSE* 8(1), 3–18 (2012), <https://doi.org/10.1007/s11334-011-0155-2>
3. Demuth, A., Riedl-Ehrenleitner, M., Nöhrer, A., Hehenberger, P., Zeman, K., Egyed, A.: Designspace: an infrastructure for multi-user/multi-tool engineering. In: Wainwright, R.L., Corchado, J.M., Bechini, A., Hong, J. (eds.) Proceedings of the 30th Annual ACM Symposium on Applied Computing, Salamanca, Spain, April 13–17, 2015. pp. 1486–1491. ACM (2015), <https://doi.org/10.1145/2695664.2695697>
4. Diskin, Z.: Model Synchronization: Mappings, Tiles, and Categories. In: Fernandes, J.M., Lämmel, R., Visser, J., Saraiva, J. (eds.) GTTSE. Lecture Notes in Computer Science, vol. 6491, pp. 92–165. Springer (2011)
5. Diskin, Z.: Compositionality of update propagation: Lax putput. In: Eramo, R., Johnson, M. (eds.) Proceedings of the 6th International Workshop on Bidirectional Transformations co-located with The European Joint Conferences on Theory and Practice of Software, BX@ETAPS 2017, Uppsala, Sweden, April 29, 2017. CEUR Workshop Proceedings, vol. 1827, pp. 74–89. CEUR-WS.org (2017), <http://ceur-ws.org/Vol-1827/paper12.pdf>
6. Diskin, Z., Gholizadeh, H., Wider, A., Czarnecki, K.: A three-dimensional taxonomy for bidirectional model synchronization. *Journal of System and Software* 111, 298–322 (2016), <https://doi.org/10.1016/j.jss.2015.06.003>
7. Diskin, Z., König, H., Lawford, M.: Multiple model synchronization with multiary delta lenses. In: Russo, A., Schürr, A. (eds.) Fundamental Approaches to Software Engineering, 2018. Lecture Notes in Computer Science, vol. 10802, pp. 21–37. Springer (2018), [https://doi.org/10.1007/978-3-319-89363-1\\_2](https://doi.org/10.1007/978-3-319-89363-1_2)
8. Diskin, Z., König, H., Lawford, M., Maibaum, T.: Toward product lines of mathematical models for software model management. In: Seidl, M., Zschaler, S.

- (eds.) *Software Technologies: Applications and Foundations - STAF 2017 Collocated Workshops*, Marburg, Germany, July 17-21, 2017, Revised Selected Papers. *Lecture Notes in Computer Science*, vol. 10748, pp. 200–216. Springer (2017), [https://doi.org/10.1007/978-3-319-74730-9\\_19](https://doi.org/10.1007/978-3-319-74730-9_19)
9. Diskin, Z., Xiong, Y., Czarnecki, K.: Specifying Overlaps of Heterogeneous Models for Global Consistency Checking. In: *MoDELS Workshops*. LNCS, vol. 6627, pp. 165–179. Springer (2010)
  10. Diskin, Z., Xiong, Y., Czarnecki, K.: From State- to Delta-Based Bidirectional Model Transformations: the Asymmetric Case. *Journal of Object Technology* 10, 6: 1–25 (2011)
  11. Diskin, Z., Xiong, Y., Czarnecki, K., Ehrig, H., Hermann, F., Orejas, F.: From state-to delta-based bidirectional model transformations: the symmetric case. In: *MODELS*, pp. 304–318. Springer (2011)
  12. Ehrig, H., Ehrig, K., Prange, U., Taenzer, G.: *Fundamentals of Algebraic Graph Transformation* (2006)
  13. Ehrig, H., Ehrig, K., U.Prange, Taentzer, G.: *Fundamentals of algebraic graph transformations*. Springer (2006)
  14. Foster, J.N., Greenwald, M.B., Moore, J.T., Pierce, B.C., Schmitt, A.: Combinators for bi-directional tree transformations: a linguistic approach to the view update problem. In: Palsberg, J., Abadi, M. (eds.) *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005*, Long Beach, California, USA, January 12-14, 2005. pp. 233–246. ACM (2005), <http://doi.acm.org/10.1145/1040305.1040325>
  15. Foster, J.N., Greenwald, M.B., Moore, J.T., Pierce, B.C., Schmitt, A.: Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 29(3), 17 (2007)
  16. Hermann, F., Ehrig, H., Ermel, C., Orejas, F.: Concurrent model synchronization with conflict resolution based on triple graph grammars. *Fundamental Approaches to Software Engineering* pp. 178–193 (2012)
  17. Hermann, F., Ehrig, H., Orejas, F., Czarnecki, K., Diskin, Z., Xiong, Y.: Correctness of model synchronization based on triple graph grammars. In: *MODELS*, pp. 668–682. Springer (2011)
  18. Hofmann, M., Pierce, B.C., Wagner, D.: Symmetric lenses. In: Ball, T., Sagiv, M. (eds.) *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011*, Austin, TX, USA, January 26-28, 2011. pp. 371–384. ACM (2011), <http://doi.acm.org/10.1145/1926385.1926428>
  19. Hofmann, M., Pierce, B.C., Wagner, D.: Edit lenses. In: Field, J., Hicks, M. (eds.) *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012*, Philadelphia, Pennsylvania, USA, January 22-28, 2012. pp. 495–508. ACM (2012), <http://doi.acm.org/10.1145/2103656.2103715>
  20. Johnson, M., Rosebrugh, R.D.: Lens put-put laws: monotonic and mixed. *ECE-ASST* 49 (2012)
  21. Johnson, M., Rosebrugh, R.D.: Unifying set-based, delta-based and edit-based lenses. In: Anjorin, A., Gibbons, J. (eds.) *Proceedings of the 5th International Workshop on Bidirectional Transformations, Bx 2016*, co-located with The European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 8, 2016. *CEUR Workshop Proceedings*, vol. 1571, pp. 1–13. CEUR-WS.org (2016), [http://ceur-ws.org/Vol-1571/paper\\_13.pdf](http://ceur-ws.org/Vol-1571/paper_13.pdf)
  22. Johnson, M., Rosebrugh, R.D.: Symmetric delta lenses and spans of asymmetric delta lenses. *Journal of Object Technology* 16(1), 2:1–32 (2017), <https://doi.org/10.5381/jot.2017.16.1.a2>
  23. Johnson, M., Rosebrugh, R.D.: Cospans and symmetric lenses. In: *Conference Companion of the 2nd International Conference on Art, Science, and Engineer-*

- ing of Programming, Nice, France, April 09-12, 2018. pp. 21–29 (2018), <http://doi.acm.org/10.1145/3191697.3191717>
24. Johnson, M., Rosebrugh, R.D., Wood, R.J.: Lenses, fibrations and universal translations. *Mathematical Structures in Computer Science* 22(1), 25–42 (2012), <https://doi.org/10.1017/S0960129511000442>
  25. Knapp, A., Mossakowski, T.: Multi-view consistency in UML. *CoRR* abs/1610.03960 (2016), <http://arxiv.org/abs/1610.03960>
  26. König, H., Diskin, Z.: Efficient consistency checking of interrelated models. In: *Modelling Foundations and Applications - 13th European Conference, ECMFA 2017, Held as Part of STAF 2017, Marburg, Germany, July 19-20, 2017, Proceedings*. pp. 161–178 (2017), [https://doi.org/10.1007/978-3-319-61482-3\\_10](https://doi.org/10.1007/978-3-319-61482-3_10)
  27. Macedo, N., Cunha, A., Pacheco, H.: Towards a framework for multidirectional model transformations. In: *Proceedings of the Workshops of the EDBT/ICDT 2014 Joint Conference (EDBT/ICDT 2014), Athens, Greece, March 28, 2014*. pp. 71–74 (2014), <http://ceur-ws.org/Vol-1133/paper-11.pdf>
  28. Marussy, K., Semeráth, O., Varró, D.: Incremental view model synchronization using partial models. In: Wasowski, A., Paige, R.F., Haugen, Ø. (eds.) *Proceedings of the 21th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems, MODELS 2018, Copenhagen, Denmark, October 14-19, 2018*. pp. 323–333. ACM (2018), <https://doi.org/10.1145/3239372.3239412>
  29. Orejas, F., Boronat, A., Ehrig, H., Hermann, F., Schölzel, H.: On propagation-based concurrent model synchronization. *ECEASST* 57 (2013), <http://journal.ub.tu-berlin.de/eceasst/article/view/871>
  30. Rabbi, F., Lamo, Y., Yu, I.C., Kristensen, L.M.: A diagrammatic approach to model completion. In: *Proceedings of the 4th Workshop on the Analysis of Model Transformations co-located with the 18th International Conference on Model Driven Engineering Languages and Systems (MODELS 2015), Ottawa, Canada, September 28, 2015*. pp. 56–65 (2015), <http://ceur-ws.org/Vol-1500/paper7.pdf>
  31. Stevens, P.: Bidirectional model transformations in qvt: semantic issues and open questions. *Software & Systems Modeling* 9(1), 7–20 (2010)
  32. Stevens, P.: Bidirectional transformations in the large. In: *20th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems, MODELS 2017, Austin, TX, USA, September 17-22, 2017*. pp. 1–11 (2017), <https://doi.org/10.1109/MODELS.2017.8>
  33. Trollmann, F., Albayrak, S.: Extending model to model transformation results from triple graph grammars to multiple models. In: *Theory and Practice of Model Transformations - 8th International Conference, ICMT 2015, Held as Part of STAF 2015, L'Aquila, Italy, July 20-21, 2015, Proceedings*. pp. 214–229 (2015), [https://doi.org/10.1007/978-3-319-21155-8\\_16](https://doi.org/10.1007/978-3-319-21155-8_16)
  34. Trollmann, F., Albayrak, S.: Extending model synchronization results from triple graph grammars to multiple models. In: *Theory and Practice of Model Transformations - 9th International Conference, ICMT 2016, Held as Part of STAF 2016, Vienna, Austria, July 4-5, 2016, Proceedings*. pp. 91–106 (2016), [https://doi.org/10.1007/978-3-319-42064-6\\_7](https://doi.org/10.1007/978-3-319-42064-6_7)
  35. Trollmann, F., Albayrak, S.: Decision points for non-determinism in concurrent model synchronization with triple graph grammars. In: *Theory and Practice of Model Transformation - 10th International Conference, ICMT 2017, Held as Part of STAF 2017, Marburg, Germany, July 17-18, 2017, Proceedings*. pp. 35–50 (2017), [https://doi.org/10.1007/978-3-319-61473-1\\_3](https://doi.org/10.1007/978-3-319-61473-1_3)