



# Avoiding unnecessary information loss: correct and efficient model synchronization based on triple graph grammars

Lars Fritsche<sup>1</sup> · Jens Kosiol<sup>2</sup> · Andy Schürr<sup>1</sup> · Gabriele Taentzer<sup>2</sup>

Published online: 8 September 2020  
© The Author(s) 2020

## Abstract

Model synchronization, i.e., the task of restoring consistency between two interrelated models after a model change, is a challenging task. Triple graph grammars (TGGs) specify model consistency by means of rules that describe how to create consistent pairs of models. These rules can be used to automatically derive further rules, which describe how to propagate changes from one model to the other or how to change one model in such a way that propagation is guaranteed to be possible. Restricting model synchronization to these derived rules, however, may lead to unnecessary deletion and recreation of model elements during change propagation. This is inefficient and may cause unnecessary information loss, i.e., when deleted elements contain information that is not represented in the second model, this information cannot be recovered easily. *Short-cut rules* have recently been developed to avoid unnecessary information loss by reusing existing model elements. In this paper, we show how to automatically derive (*short-cut*) *repair rules* from short-cut rules to propagate changes such that information loss is avoided and model synchronization is accelerated. The key ingredients of our rule-based model synchronization process are these repair rules and an *incremental pattern matcher* informing about suitable applications of them. We prove the termination and the correctness of this synchronization process and discuss its completeness. As a proof of concept, we have implemented this synchronization process in eMoflon, a state-of-the-art model transformation tool with inherent support of bidirectionality. Our evaluation shows that repair processes based on (*short-cut*) *repair rules* have considerably decreased information loss and improved performance compared to former model synchronization processes based on TGGs.

**Keywords** Bidirectional transformation · Model synchronization · Triple graph grammar · Incremental pattern matching · Change propagation

## 1 Introduction

The close collaboration of multiple disciplines such as electrical engineering, mechanical engineering, and software

engineering in system design often leads to discipline-spanning system models [27]. Keeping models synchronized by checking and preserving their consistency can be a challenging problem which is not only subject to ongoing research but also of practical interest for industrial applications. Model-based engineering has become an important technique to cope with the increasing complexity of modern software systems. Various *bidirectional transformation* (bx) approaches [3,14] for models have been suggested to deal with model (view) synchronization and consistency. Across these different approaches the following are important research topics [13,15,26,31–33,47]: *incrementality*, i.e., achieving runtime/complexity dependent on the size of the model change, not on the model size, and *least change*, i.e., keeping the resulting model as similar as possible to the original one while restoring consistency. In this work, we extend

---

✉ Lars Fritsche  
lars.fritsche@es.tu-darmstadt.de

✉ Jens Kosiol  
kosiolje@mathematik.uni-marburg.de

Andy Schürr  
andy.schuerr@es.tu-darmstadt.de

Gabriele Taentzer  
taentzer@mathematik.uni-marburg.de

<sup>1</sup> Real-Time Systems Lab, TU Darmstadt, Darmstadt, Germany

<sup>2</sup> Faculty of Mathematics and Computer Science, Philipps-Universität Marburg, Marburg, Germany

synchronization approaches based on *triple graph grammars* by specific *repair rules* to increase incrementality and efficiency and to decrease the amount of change that occurs during synchronization. We show how to avoid unnecessary information loss in model synchronization for scenarios *in which one model is changed at a time*. Throughout this paper we stick to this scenario of model synchronization.

The more general case of *concurrent model synchronization* where both models have been altered is left to future work.

Triple Graph Grammars (TGGs) [51] are a declarative, rule-based bidirectional transformation approach, which allows synchronizing models of two different views (usually called the *source* and *target domain* in the TGG-related literature). The purpose of a TGG is to define a consistency relationship between pairs of models in a rule-based manner by defining traces between their elements. Given a TGG, its rules can be automatically *operationalized* into *source* and *forward rules*. While the source rules are used to build up models of the source domain, forward rules translate them to the target domain and thereby establish traces between corresponding model elements. Analogously, target models can be propagated to the source domain by using *target* and *backward rules* that can be automatically deduced as well. To avoid redundancy in our presentation, we stick to forward propagation throughout this paper.

In [51], a simple batch-oriented synchronization process was presented, which just re-translates the whole source model after each change using forward rules. Several incremental synchronization processes based on TGGs have been presented in the literature thereafter. A process is considered to be incremental if the target model is not recomputed from scratch but unaffected model parts are preserved as much as possible.<sup>1</sup> To obtain an incremental synchronization process, two basic strategies have been pursued (in combinations): (i) The synchronization algorithm takes additional information of forward rules into account. This information might consist of precedence relations over rules [40], dependency information on model elements w.r.t. their creation [26,50], a maximal, still consistent submodel [30], or information about broken matches of forward rules provided by an incremental pattern matcher [41,42]. (ii) The actual propagation of changes in a synchronization process is not based on the application of forward rules exclusively but also uses additional rules. To propagate a deletion on the source part, almost all approaches support to revoke an application of a forward rule. The revocation of rule applications is formalized as inverse rule applications in, e.g., [40]. Also, custom-

made rules have been used in synchronization algorithms that describe specific kinds of model edits in any modeling language [24] or in a concrete modeling language [10]. Moreover, generalized forward rules have been defined which allow for re-use of elements [24,27,50]. Summarizing, several approaches for incremental model synchronization based on TGGs have been presented in the literature. Some of them such as [26,27] are informally presented without any guarantee to reestablish the consistency of modified models. Others present their synchronization approaches formally and show their correctness but are only applicable under restricted circumstances [30] or have not been implemented yet, such as [50]. Hence, *we still miss a TGG-based model synchronization approach that avoids unnecessary information loss, is proven to be correct, and is efficiently implemented*.

In this article, we present an incremental model synchronization approach based on an extended set of TGG rules. In [22], we introduced short-cut rules for handling complex consistency-preserving model updates while avoiding unnecessary information loss. A short-cut rule replaces one rule application with another one while preserving involved model elements (instead of deleting and re-creating them). We deduce source and forward rules from short-cut rules to support complex model edits and their synchronization with the target domain. We present an incremental model synchronization algorithm based on short-cut rules and show its correctness.

We implemented our synchronization approach in eMoflon [43,57,58], a state-of-the-art bidirectional model transformation tool, and evaluate it. Being based on eMoflon, we are able to extend the synchronization process suggested by Leblebici et al. [41,42] and rely on information provided by an incremental pattern matcher also to detect when and where to apply our derived repair rules. However, the construction and derivation of these are general and could extend other suggested TGG-based synchronization processes as well. The results of our evaluation show that, compared to model synchronization in eMoflon without short-cut repair rules, the application of these repair rules allows reacting to model changes in a less invasive way by preserving information. Besides, it shows more efficiency.

This paper extends the work in [23]. Beyond [23], we

- present the actual synchronization process in pseudocode and prove its correctness and termination (based on the results obtained in [23,41,42]),
- extend our approach to deal with *filter NACs* (a specific kind of negative application conditions in forward rules),
- describe the implementation, especially the tool architecture, in more detail,
- extend the evaluation by investigating the expressiveness of short-cut repair rules at the practical example of code refactorings [21], and

<sup>1</sup> Ideally, the runtime (complexity) of a synchronization should depend on the size of the change to the source model and not on the sizes of the source and the target model [26]. This requirement is a good motivation for incremental synchronization.

- consider the related work more comprehensively.

The rest of this paper is organized as follows. In Sect. 2, we give an informal overview of our model synchronization approach. It shall allow readers to grasp the general idea without working through the technical details. In Sect. 3, we recall triple graph grammars. The construction of short-cut rules and their properties are presented in Sect. 4, while Sect. 5 introduces the derivation of repair rules. Section 6 focuses on the implemented synchronization algorithm and its formal properties. To be understandable to readers who are not experts on algebraic graph transformation, we use a set-theoretical notion in these more technical sections, in contrast to the original contribution in [23] which is based on category theory. Section 7 describes the implementation of our model synchronization algorithm in eMoflon, focussing on the tool architecture. Our synchronization approach is evaluated in Sect. 8. Finally, we discuss related work in Sect. 9 and conclude with pointers to future work in Sect. 10. "Appendix" presents the rule set used for our evaluation.

## 2 Informal introduction to TGG-based model synchronization

In this section, we illustrate our approach to model synchronization. Using a simple example, we will explain the basic concepts as well as all main ingredients for our new synchronization process. Reading this section and having a passing view on the synchronization algorithm (Sect. 6.2), evaluation (Sect. 8), and related works (Sect. 9) should give an adequate impression of the core ideas of our work.

Graph transformations, and triple graph grammars in particular, are a suitable formal framework to reason about and to implement model transformations and synchronizations [9,18].<sup>2</sup> A triple graph consists of three graphs, namely the *source*, *target*, and *correspondence graph*. The latter encodes which elements of source and target graph *correlate to each other*. This is done by mapping each element of the correspondence graph to an element of the source graph as well as to an element of the target graph (formally these are two *graph morphisms*). Elements connected via such a mapping are considered to be correlated.

*Triple graph grammars* (TGGs) [51] declaratively define how consistent models co-evolve. This means that a triple graph is considered to be consistent if it can be derived from a start triple (e.g., the empty graph) using the rules of the given grammar. Furthermore, the rules can automatically be

<sup>2</sup> Therefore, we will use the terms "graph" and "model" interchangeably in this paper. We will stick to the graph terminology in more formal sections.




Fig. 1 Example: type graph

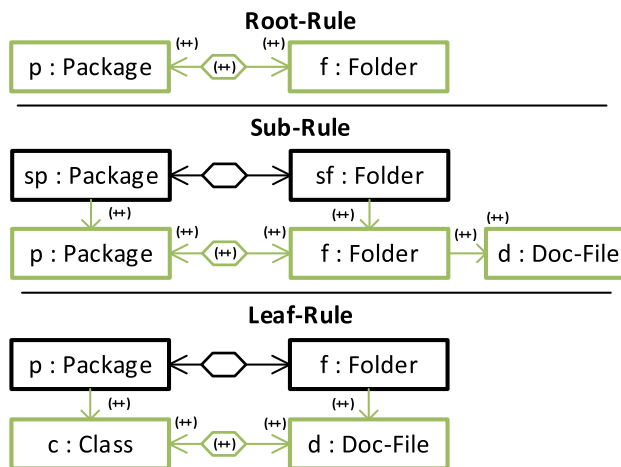


Fig. 2 Example: TGG rules

*operationalized* to obtain new kinds of rules, e.g., for translation/synchronization processes.

We illustrate our model synchronization process by synchronizing a Java AST (abstract syntax tree) and a custom documentation model as example. This example has been basically introduced by Leblebici et al. [44]; it is slightly modified to demonstrate the core concepts of our approach. Note, however, that the evaluation in Sect. 8 is based on a larger and more complex TGG consisting of 24 rules (as presented in "Appendix").

For model synchronization, we consider a Java AST model as *source* model and its documentation model as *target* model, i.e., changes in a Java AST model have to be transferred to its documentation model and vice versa. Note that we do not consider concurrent model synchronization, i.e., concurrent changes to both sides that have to be synchronized. Figure 1 depicts the type graph that describes the syntax of our example triple graphs. It shows a *Package* hierarchy and *Classes* as the source side, a *Folder* hierarchy with *Doc-Files* as target side and correspondence types in between depicted as hexagons. Furthermore, *Doc-Files* have an attribute *content* which is of type *String*. Note that, in our example, there are two correspondence types which

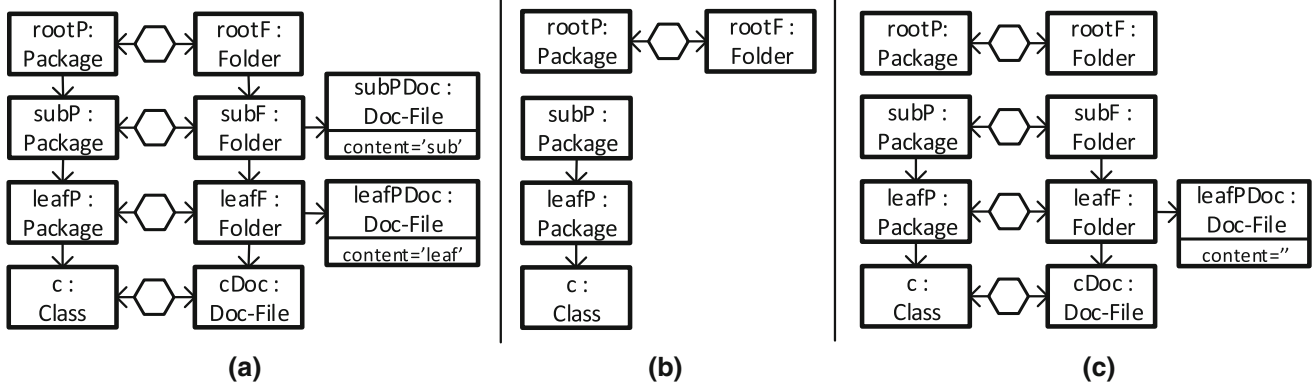


Fig. 3 Exemplary synchronization scenario

can be distinguished by the type of elements they connect on both sides.

*TGG rules* Figure 2 shows the rule set of our example TGG consisting of three rules (assuming an empty start graph):

*Root-Rule* creates a root *Package* together with a root *Folder* and a correspondence link in between. This rule has an empty precondition and creates elements only; they are depicted in green and with the annotation  $(++)$ . *Sub-Rule* creates a *Package* and *Folder* hierarchy given that an already correlated *Package* and *Folder* pair exists. Finally, *Leaf-Rule* creates a *Class* and a *Doc-File* under the same precondition as *Sub-Rule*.

TGG rules can be used to generate triple graphs; triple graphs generated by them are consistent by definition. An example is depicted in Fig. 3(a) which can be generated by first applying *Root-Rule* followed by two applications of *Sub-Rule* and an application of *Leaf-Rule*: Starting with the empty triple graph, the first rule application just creates the elements  $rootP$  and  $rootF$  and the correspondence element in between. The second rule application matches these elements and creates  $subP$ ,  $subF$ ,  $subPDoc$ , their respective incoming edges, and the correspondence element between  $subP$  and  $subF$ . The other two rule applications are performed similarly.

*Operationalization of TGG rules* A TGG can also be used for *translating* a model of one domain to a correlated model of a second domain. Moreover, a TGG offers support for *model synchronization*, i.e., for restoring the consistency of a triple graph that has been altered on one side. For these purposes, each TGG rule has to be operationalized to two kinds of rules: A *source* rules enable changes of source models (e.g., as performed by a user) while forward rules translate such changes to the target model.<sup>3</sup> The result of applying a source rule followed by an application of its corresponding forward rule yields the same result as applying the TGG rule they

<sup>3</sup> Analogously, *target* and *backward* rules can be derived.

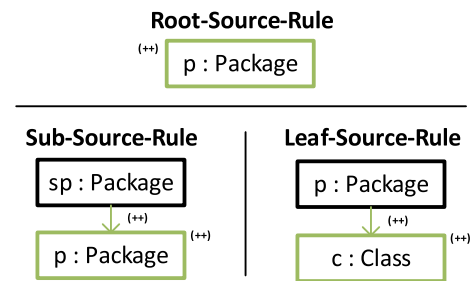


Fig. 4 Example: TGG source rules

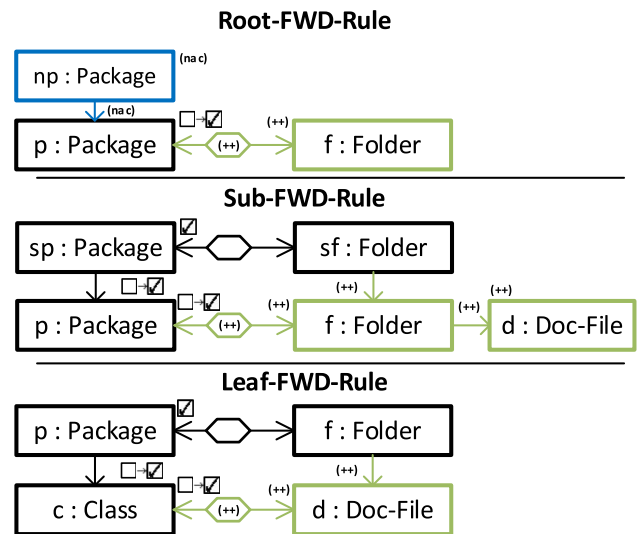


Fig. 5 Example: TGG forwardRules

originate from. Figure 4 shows the resulting *source* rules for our example TGG.

*Forward translation rules* Figure 5 depicts the resulting *forward* rules. They have a similar structure compared to their original TGG rules with three important differences. First, elements on the source side are now considered as context and as such have to be matched as a precondition for this rule to be applicable. Second, since we consider elements

on the source side to already be present, we have to mark whether an element has already been translated or not. A  $\checkmark$  annotation can be found on source elements which must have been translated before. On the other hand,  $\square \rightarrow \checkmark$  annotations indicate that applying this rule would mark this element as translated. This annotation can be found at elements that are created by the original TGG rule. Possible formalizations of these marking are given, e.g., in [29,42]. The third difference is the use of negative application conditions (NACs) [17] which are indicated with a (nac) and depicted in blue. Using NACs, we are able to not only define necessary structure that has to be found but also the explicit absence of structural elements as in *Root-FWD-Rule* where we forbid `subP` to have a parent package. The theory behind these so-called *filter NACs* is formalized by Hermann et al. [29] and they can be derived automatically from the rules of a given TGG when computing its forward rules.

Using these rules, we can translate Java AST to documentation models. Considering the one on the source side of the triple graph in Fig. 3(a), it is translated to a documentation model such that the result is the complete graph depicted in this part of the figure. To obtain this result we apply *Root-FWD-Rule* at the root *Package*, *Sub-FWD-Rule* at *Packages* `subP` and `leafP`, and finally *Leaf-FWD-Rule* at *Class* `c`. Note that *Sub-FWD-Rule*, for example, is applicable when matching *Packages* `sp` and `p` of the rule to the *Packages* `rootP` and `subP` of the source graph, respectively, since `rootP` was marked as translated by the application of *Root-FWD-Rule*.

Without the NAC in *Root-FWD-Rule*, this rule would also be applicable at the elements `subP` and `leafP`. Applying *Root-FWD-Rule* and translating these elements with it, however, would result in the edges from their parent *Packages* not being translatable any longer: There is no rule in our TGG rule set that creates edges between packages only. Hence, NACs can direct the translation process to avoid these dead-ends. Filter NACs are derived such that they prevent rule applications leading to dead-ends, only.

*Existing approaches to model synchronization* Given a triple graph such as the one in Fig. 3(a), a developer may want to split the modeled project into multiple ones. For this purpose, a subpackage such as `subP` shall become a root package. Since `subP` was created and translated as a sub package rather than a root element, this model change introduces an inconsistency. To resolve this issue, the approaches presented in [26,40–42] and, to a certain degree, also the one in [30] revert the translation of `subP` into `subF` and re-translate `subP` with an appropriate translation rule such as *Root-FWD-Rule*. Reverting the former translation step may lead to further inconsistencies as we remove elements that were needed as context elements by other applications of forward rules. The result is a reversion of all translation steps except for the first one which translates the original

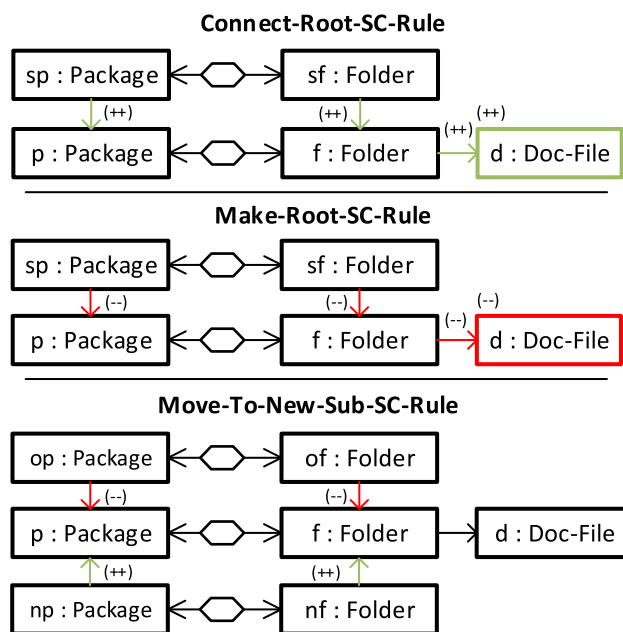


Fig. 6 Short-cut rules

root element. The result is shown in Fig. 3(b). Thereafter, the untranslated elements can be re-translated yielding the result graph in (c). This example shows that this synchronization approach may delete and re-create a lot of similar structures which appears to be inefficient. Second, it may lose information that exists on the target side only, e.g., documentation saved in the content attribute which is empty now as it cannot be restored from the source side only. Such an information loss is unnecessary as we will show below. Instead of deleting elements and recreating them, we will present a synchronization process that aims to preserve information as much as possible.

*Model synchronization with short-cut repair* In [22], we introduce short-cut rules as a kind of sequential rule composition mechanism that allows replacing one rule application with another one while elements are preserved (instead of deleted and recreated).

Figure 6 depicts three short-cut rules which can be derived from our original three TGG rules. The first two, *Connect-Root-SC-Rule* and *Make-Root-SC-Rule*, are derived from *Root-Rule* and *Sub-Rule*. The upper short-cut rule replaces an application of *Root-Rule* with one of *Sub-Rule* and turns root elements into sub elements. In contrast, the lower short-cut rule replaces an application of *Sub-Rule* with one of *Root-Rule*, thus, turning sub elements into root elements. Both short-cut rules preserve the model elements present in their corresponding TGG rules and solely create elements that do not exist yet (++) , or delete those depicted in red and annotated with (--) which became superfluous. The third short-cut rule *Move-To-New-Sub-SC-Rule* relocates sub ele-



ments and replaces a *Sub-Rule* application with another one of the same kind.

A short-cut rule is constructed by overlapping two rules with each other where the first one is the replaced and the second the replacing rule. Overlapped elements are preserved such as *p* and *f* in *Connect-Root-SC-Rule*. Created elements that are not overlapped fall into two categories. If the element was created in the replaced rule but is superfluous in the replacing rule, it is deleted, e.g., *d* in *Make-Root-SC-Rule*. On the other hand, if the element was not created by the replaced rule but by the replacing rule, then the element is created, e.g., *d* in *Connect-Root-SC-Rule*. Context elements can be mapped as well while unmapped context elements from both rules are glued onto the final short-cut rule, e.g., *op* and *of* which are context in the replaced rule, and *np* and *nf* which are context in the replacing rule. Since there are many possible overlaps for each pair of rules, constructing a reasonable set of short-cut rules depends on the concrete example TGG and the requirement for advanced model changes that go beyond the standard capabilities of TGG-based model synchronizers. Usually, it is worthwhile to construct short-cut rules for frequent model changes in order to increase the synchronization efficiency and decrease information loss in these cases.

In our example above, the user wants to transform the triple graph in Fig. 3(a) to the one in (c). Using *Make-Root-SC-Rule* and matching the *Packages* *sp* and *p* to the *Packages* *rootP* and *subP* of the model (a) (and the correspondence nodes and *Folder* accordingly), this transformation is performed with a single rule application. Analogously, the triple graph (c) can be directly transformed backwards to (a) using *Connect-Root-SC-Rule*. Thus, these rules allow for complex user-edits on both, source and target side; they preserve the consistency of the model. However, there are also scenarios where applying a short-cut rule may lead to an inconsistent state of the resulting triple graph. A simple example is that of applying *Connect-Root-SC-Rule* in order to connect *subP* and *subF* with *rootP* and *rootF*, respectively. The result would be a cycle in both, the *Package* and the *Folder* hierarchies; this model is no longer in the language of our example TGG. In Sect. 4, we present sufficient conditions for the application of short-cut rules to avoid such cases.

*Operationalization of short-cut rules* Short-cut rules transform both models at once as TGG rules usually do and therefore, they cannot cope with the change of a single model. Hence, similar to TGG rules, we have to operationalize them, thereby obtaining *short-cut source* and *short-cut repair* rules. Figure 7 depicts the *short-cut source* rules which are derived analogously to those of standard TGG rules. In order to be able to handle the deleted edge between *rootP* and *subP*, as deleted by *Make-Root-Source-Rule*, for example, a repair rule is needed that adapts the target graph accordingly by deleting the now superfluous edge between *rootF*

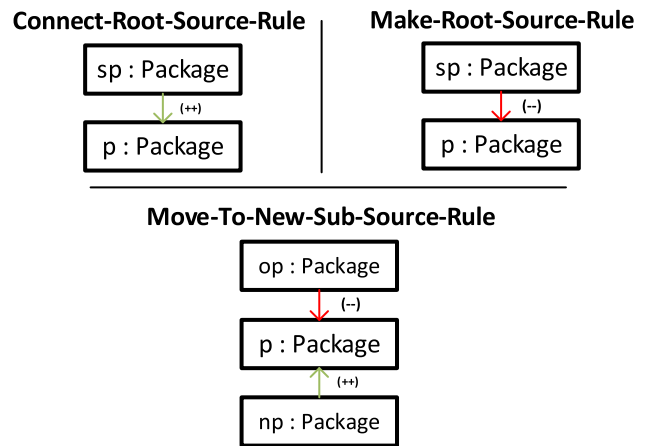


Fig. 7 Short-cut source rules

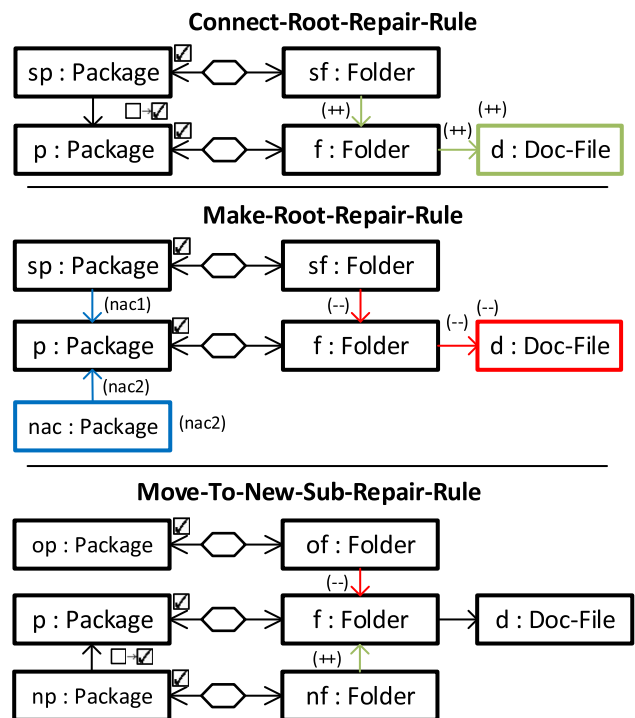


Fig. 8 Repair rules

and *subF*. Figure 8 depicts the resulting repair rules derived from the short-cut rules in Fig. 6. A short-cut rule is forward operationalized by removing deleted elements from the rule's source graph since these deletions have already happened. Furthermore, created source elements become context because we expect them to already exist, e.g., through the a prior source rule application. Finally, since short-cut rules transform an application of one rule into that of another, filter NACs are added during operationalization to comply with application conditions of the replacing rule which naturally have to hold when applying the short-cut rule. Hence, *Make-Root-Repair-Rule* is only applicable and can turn *subF* into

a root *Folder* if *subP* has no parent packages and, thus, is indeed a root *Package* itself. Note that *Root-FWD-Rule* is only applicable if *subP* has no parent packages, which *Make-Root-Repair-Rule* has to incorporate as well. For this reason, *Make-Root-Repair-Rule* contains *nac1*, which forbids *rootP* to be the parent package of *subP* and *nac2*, which forbids *subP* to have any other parent packages than *rootP*.

*Short-cut repair* rules allow propagating graph changes directly to the other graph to restore consistency. Revisiting our example of Fig. 3, we are now able to use *Make-Root-Repair-Rule* to propagate the deleted edge between *subP* and *rootP* by deleting the corresponding edge between *subF* and *rootF* and the now superfluous *Doc-File subPDoc*. The result is the consistent triple graph again depicted in Fig. 3(c) with the content attribute of *leafPDoc* containing the value ‘leaf’. So, this repair does not cause information loss and allows skipping the costly reversion process with the intermediate result in Fig. 3(b).

Summarizing, the user edit of removing the edge between *Packages rootP* and *subP* corresponds to the source rule of *Make-Root-SC-Rule*, namely *Make-Root-Source-Rule*, and the according update to the target side is performed by *Make-Root-Repair-Rule* which is the corresponding repair rule. Together, they perform an edit step structurally equivalent to the one depicted by the triple graphs in Fig. 3(a), (c); however, the value of the attribute content does not get lost. Alternatively, this step can be obtained by applying the short-cut rule *Make-Root-SC-Rule*. This is not a coincidence: In [23, Theorem 7], we showed that applying the source rule of a short-cut rule (which corresponds to a user edit on the source part only) followed by an application of the corresponding repair rule at the according match is the same as applying the original short-cut rule.

### 3 Preliminaries: triple graphs, triple graph grammars and their operationalizations

In this section, we recall triple graph grammars (TGGs) and their operationalization [51]. Our derivation of repair rules is based on the construction of so-called short-cut rules [22], which we recall as well. For simplicity, we stick with set-theoretic definitions of the involved concepts (in contrast to category-theoretic ones as, e.g., in [17,18,22,23]). Moreover, while we provide formal definitions for central notions, we will just explain others and provide references for their formal definitions.

#### 3.1 Graphs, triple graphs, and their transformations

*Graphs* and their (rule-based) *transformations* are suitable to formalize various kinds of models and their evolution, in

particular of EMF models [9].<sup>4</sup> In the context of this work, a *graph* consists of a set of nodes and a set of directed edges which connect nodes. Graphs may be related by *graph morphisms*, and a *triple graph* consists of three graphs connected by two graph morphisms.

**Definition 1** (*Graph, graph morphism, triple graph, and triple graph morphism*) A *graph*  $G = (V, E, s, t)$  consists of a set  $V$  of vertices, a set  $E$  of edges, and source and target functions  $s, t : E \rightarrow V$ . An *element*  $x$  of  $G$  is a node or an edge, i.e.,  $x \in V$  or  $x \in E$ . A *graph morphism*  $f : G \rightarrow H$  between graphs  $G = (V_G, E_G, s_G, t_G)$  and  $H = (V_H, E_H, s_H, t_H)$  consists of two functions  $f_V : V_G \rightarrow V_H$  and  $f_E : E_G \rightarrow E_H$  that are compatible with the assignment of source and target to edges, i.e.,  $f_V \circ s_G = s_H \circ f_E$  and  $f_V \circ t_G = t_H \circ f_E$ . Given a fixed graph  $TG$ , a *graph typed over TG* is a graph  $G$  together with a graph morphism  $type_G : G \rightarrow TG$ . A *typed graph morphism*  $f : (G, type_G) \rightarrow (H, type_H)$  between typed graphs is a graph morphism  $f : G \rightarrow H$  that respects the typing, i.e.,  $type_G = type_H \circ f$  (componentwise). A (typed) graph morphism  $f = (f_V, f_E)$  is *injective* if both  $f_V$  and  $f_E$  are.

A *triple graph*  $G = (G_S \xleftarrow{\sigma_G} G_C \xrightarrow{\tau_G} G_T)$  consists of three graphs  $G_S, G_C, G_T$ , called *source, correspondence, and target graph*, and two graph morphisms  $\sigma_G : G_C \rightarrow G_S$  and  $\tau_G : G_C \rightarrow G_T$ , called *source and target correspondence morphism*. A *triple graph morphism*  $f : G \rightarrow H$  between two triple graphs  $G$  and  $H$  consists of three graph morphisms  $f_S : G_S \rightarrow H_S, f_C : G_C \rightarrow H_C$  and  $f_T : G_T \rightarrow H_T$  such that  $\sigma_H \circ f_C = f_S \circ \sigma_G$  and  $\tau_H \circ f_C = f_T \circ \tau_G$ . Given a fixed triple graph  $TG$ , a *triple graph typed over TG* is a triple graph  $G$  together with a triple graph morphism  $type_G : G \rightarrow TG$ . Again, *typed triple graph morphisms* are triple graph morphisms that respect the typing. A (typed) triple graph morphism  $f = (f_S, f_C, f_T)$  is *injective* if  $f_S, f_C$ , and  $f_T$  all are.

**Example 1** Figure 3 depicts three triple graphs; their common type graph is depicted in Fig. 1. The typing morphism is indicated by annotating the elements of the triple graphs with the types to which they are mapped in the type graph. The nodes in the triple graphs are of types *Package, Folder, Class, and Doc-File*. In each case, the source graph is depicted to the left and the target graph to the right. The hexagons in the middle constitute the correspondence graphs. Formally, the edges from the correspondence graphs to source and target graphs are morphisms: The edges encode how an individual correspondence node is mapped by the correspondence morphisms. For example, the nodes *rootP* and *rootF* of types *Package* and *Folder* correspond to each other as they

<sup>4</sup> Therefore in this paper, we use the terms *graph* and *model* interchangeably. In the formal parts, we will consequently speak of graphs following the formal literature.

share the same correspondence node as preimage under the correspondence morphisms.

Rules offer a declarative means to specify transformations of (triple) graphs. While classically rewriting of triple graphs has been performed using non-deleting rules only, we define a less restricted notion of rules<sup>5</sup> right away since short-cut rules and repair rules derived from them are both potentially deleting. A rule  $p$  consists of three triple graphs, namely a *left-hand side* (LHS)  $L$  and a *right-hand side* (RHS)  $R$  and an *interface*  $K$  between them. Applying such a rule to a triple graph  $G$  means to choose an injective morphism  $m$  from  $L$  to  $G$ . The elements from  $m(L \setminus l(K))$  are to be deleted; if this results in a triple graph again, the morphism  $m$  is called a *match* and  $p$  is applicable at that match. After this deletion, the elements from  $R \setminus r(K)$  are added; the whole process of applying a rule is also called a *transformation (step)*.

**Definition 2** (*Rule, transformation (step)*) A rule  $p = (L \xleftarrow{l} K \xrightarrow{r} R)$  consists of three triple graphs,  $L$ ,  $R$ , and  $K$ , called the *left-hand side*, *right-hand side*, and *interface*, respectively, and two injective triple graph morphisms  $l : K \rightarrow L$  and  $r : K \rightarrow R$ . A rule is called *monotonic*, or *non-deleting*, if  $l$  is an isomorphism. In this case we denote the rule as  $r : L \rightarrow R$ . The *inverse rule* of a rule  $p$  is the rule  $p^{-1} = (R \xleftarrow{r} K \xrightarrow{l} L)$ .

Given a triple graph  $G$ , a rule  $p = (L \xleftarrow{l} K \xrightarrow{r} R)$ , and an injective triple graph morphism  $m : L \rightarrow G$ , the rule  $p$  is *applicable at  $m$*  if

$$D := G \setminus (m(L \setminus l(K))) ,$$

is a triple graph again. Operator  $\setminus$  is understood as node- and edge-wise set-theoretic difference. The source and target functions of  $D$  are restricted accordingly. If  $D$  is a triple graph,

$$H := D \cup n(R \setminus r(K)) ,$$

is computed. Operator  $\cup$  is understood as node- and edge-wise set-theoretic union.  $n(R \setminus r(K))$  is a new copy of newly created elements.  $n$  can be extended to  $R$  by  $n(r(K)) = m(l(K))$ . The values of the source and target functions for edges from  $n(R \setminus r(K))$  with source or target node in  $K$  are determined by  $m \circ l$ , i.e.,

$$s_H(e) := m(l(r^{-1}(s_R(e))))$$

$$t_H(e) := m(l(r^{-1}(t_R(e))))$$

<sup>5</sup> As used in double pushout rewriting of graphs or objects of other adhesive categories more generally [17,39].

for such edges  $e \in n(E_R)$  with  $s_R(e) \in r_V(V_K)$  or  $t_R(e) \in r_V(V_K)$ . The whole computation is called a *transformation (step)*, denoted as  $G \Rightarrow_{p,m} H$  or just  $G \Rightarrow H$ ,  $m$  is called a *match*,  $n$  is called a *comatch* and  $D$  is the *context triple graph* of the transformation.

An equivalent definition based on computing two *pushouts*, a notion from category theory generalizing the union of sets along a common subset, serves as basis when developing a formal theory [17]. In the following and in our examples, we always assume  $K$  to be a common subgraph of  $L$  and  $R$  and the injective morphisms  $l$  and  $r$  to be the corresponding inclusions; this significantly eases the used notation. When we talk about the union of two graphs  $G_1$  and  $G_2$  along a common subgraph  $S$ , we assume that  $G_1 \cap G_2 = S$ .

To enhance expressiveness, a rule may contain *negative application conditions* (NACs) [17]. A NAC extends the LHS of a rule with a forbidden pattern: A rule is allowed to be applied only at matches which cannot be extended to any pattern forbidden by one of its NACs. If we want to stress that a rule is not equipped with NACs, we call it a *plain rule*.

**Definition 3** (*Negative application conditions*) Given a rule  $p = (L \leftarrow K \rightarrow R)$ , a set of *negative application conditions* (NACs) for  $p$  is a finite set of graphs  $NAC = \{N_1, \dots, N_k\}$  such that  $L$  is a subgraph of every one of them, i.e.,  $L \subset N_i$  for  $1 \leq i \leq k$ .

A rule  $(p = (L \leftarrow K \rightarrow R), NAC)$  with NACs is applicable at a match  $m : L \rightarrow G$  if the plain rule  $p$  is and, moreover, for none of the NACs  $N_i$  there exists an injective morphism  $x_i : N_i \rightarrow G$  such that  $x_i \circ \iota_i = m$  where  $\iota_i : L \hookrightarrow N_i$  is the inclusion of  $L$  into  $N_i$ .

**Example 2** Different sets of triple rules are depicted in Figs. 2, 5, 6, and 8. All rules in these figures are presented in an *integrated* form: Instead of displaying LHS, RHS, and the interface as three separate graphs, just one graph is presented where the different roles of the elements are displayed using markings (and color). The unmarked (black) elements constitute the interface of the rule, i.e., the context that has to be present to apply a rule. Unmarked elements and elements marked with  $(--)$  (black and red elements) form the LHS while unmarked elements and elements marked with  $(++)$  (black and green elements) constitute the RHS. Elements marked with (nac) (blue elements) extend the LHS to a NAC; different NACs for the same rule are distinguished using names.

As triple rules are depicted, their LHSs and RHSs are triple graphs themselves. For example, the LHS  $L$  of *Sub-Rule* (Fig. 2) consists of the nodes  $\text{sp}$  and  $\text{sf}$  of types *Package* and *Folder* and the correspondence node in between.

While, e.g., all rules in Fig. 2 are monotonic, *Make-Root-SC-Rule* is not as it deletes edges and a *Doc-File*. Applying *Make-Root-SC-Rule* to the triple graph (a) in Fig. 3 leads



to the triple graph (c), when *Package*-nodes  $\text{sp}$  and  $\text{p}$  (of the rule) are matched to  $\text{rootP}$  and  $\text{subP}$  (in the graph), respectively. (The *Folder* on the target part are mapped accordingly.) The rules *Connect-Root-SC-Rule* and *Make-Root-SC-Rule* are inverse to each other.

Finally, *Root-FWD-Rule* (Fig. 5) depicts a rule that is equipped with a NAC: It is applicable only at *Packages* that are not referenced by other *Packages*. This means that it is applicable at node  $\text{subP}$  in the triple graph (b) depicted in Fig. 3, but not at node  $\text{leafP}$ .

### 3.2 Triple graph grammars and their operationalization

Sets of triple graph rules can be used to define languages.

**Definition 4 (Triple graph grammar)** A triple graph grammar (TGG)  $GG = (\mathcal{R}, S)$  consists of a set of plain, monotonic triple rules  $\mathcal{R}$  and a start triple graph  $S$ . In case of typing, all rules of  $\mathcal{R}$  and  $S$  are typed over the same triple graph.

The language of a TGG  $GG$ , denoted as  $\mathcal{L}(GG)$ , is the reflexive and transitive closure of the relation induced by transformation steps via rules from  $\mathcal{R}$ , i.e.,

$$\mathcal{L}(GG) := \{H \mid S \Rightarrow_{\mathcal{R}}^* H\}$$

where  $\Rightarrow_{\mathcal{R}}^*$  denotes a finite sequence of transformation steps where each rule stems from  $\mathcal{R}$ .

The *projection of the language of a TGG to its source part* is the set

$$\mathcal{L}_S(GG) := \{G_S \mid G = (G_S \leftarrow G_C \rightarrow G_T) \in \mathcal{L}(GG)\},$$

i.e., it consists of the source graphs of the triple graphs of  $\mathcal{L}(GG)$ .

In applications, quite frequently, the start triple graph of a TGG is just the empty triple graph. We use  $\emptyset$  to denote the empty graph, the empty triple graph, and morphisms starting from the empty (triple) graph; it will always be clear from the context what is meant. To enhance expressiveness of TGGs, their rules can be extended with NACs or with some attribution concept for the elements of generated triple graphs. A recent overview of such concepts and their expressiveness can be found in [59]. In the following, we first restrict ourselves to TGGs that contain plain rules only and discuss extensions of our approach subsequently.

**Example 3** The rule set depicted in Fig. 2, together with the empty triple graph as start graph, constitutes a TGG. The triple graphs (a) and (c) in Fig. 3 are elements of the language defined by that grammar while the triple graph (b) is not.

The operationalization of triple graph rules into *source* and *forward* (or, analogously, into *target* and *backward*) rules is central to working with TGGs. Given a rule, its source rule performs the rule's actions on the source graph only while its forward rule propagates these to correspondence and target graph. This means that, for example, source rules can be used to generate the source graph of a triple graph while forward rules are then used to translate the source graph to correspondence and target side such that the result is a triple graph in the language of the TGG. Classically, this operationalization is defined for monotonic rules only [51]. We will later explain how to extend it to arbitrary triple rules. We also recall the notion of marking [41] and *consistency patterns* which can be used to check if a triple graph belongs to a given TGG.

**Definition 5 (Source and forward rule. Consistency pattern)** Given a plain, monotonic triple rule  $r = L \rightarrow R$  with  $r = (r_S, r_C, r_T)$ ,  $L = (L_S \xleftarrow{\sigma_L} L_C \xrightarrow{\tau_L} L_T)$  and  $R = (R_S \xleftarrow{\sigma_R} R_C \xrightarrow{\tau_R} R_T)$ , its *source rule* is defined as

$$r^S := (L_S \leftarrow \emptyset \rightarrow \emptyset) \xrightarrow{(r_S, id_{\emptyset}, id_{\emptyset})} (R_S \leftarrow \emptyset \rightarrow \emptyset).$$

Its *forward rule* is defined as

$$r^F := (R_S \xleftarrow{\sigma_R \circ r_C} L_C \xrightarrow{\tau_L} L_T) \xrightarrow{(id_{R_S}, r_C, r_T)} (R_S \xleftarrow{\sigma_R} R_C \xrightarrow{\tau_R} R_T).$$

We denote the left- and right-hand sides of source and forward rules of a rule  $r$  by  $L^S$ ,  $L^F$ ,  $R^S$ , and  $R^F$ , respectively.

The *consistency pattern* derived from  $r$  is the rule

$$r^C := (R_S \xleftarrow{\sigma_R} R_C \xrightarrow{\tau_R} R_T) \xrightarrow{(id_{R_S}, id_{R_C}, id_{R_T})} (R_S \xleftarrow{\sigma_R} R_C \xrightarrow{\tau_R} R_T)$$

that, upon application, just checks for the existence of the RHS of the rule without changing the instance it is applied to.

Given a rule  $r$ , each element  $x \in R_S \setminus L_S$  is called a *source marking element* of the forward rule  $r^F$ ; each element of  $L_S$  is called *required*. Given an application  $G \Rightarrow_{r^F, m^F} H$  of a forward rule  $r^F$ , the elements of  $G_S$  that have been matched by source marking elements of  $r^F$ , i.e., the elements of the set  $m^F(R_S \setminus L_S)$  are called *marked elements*.

A transformation sequence

$$G_0 \Rightarrow_{m_1^F, r_1^F} G_1 \Rightarrow_{m_2^F, r_2^F} \dots \Rightarrow_{m_t^F, r_t^F} G_t \tag{1}$$

is called *creation preserving* if no two rule applications in sequence (1) mark the same element. It is called *context preserving* if, for each rule application in sequence (1), the required elements have been marked by a previous rule application in sequence (1). If these two properties hold for sequence (1), it is called *consistently marking*. It is called *entirely marking* if every element of the common source

graph  $G_S$  of the triple graphs of this sequence is marked by a rule application in sequence (1).

The most important formal property of this operationalization is that applying a (sequence of) source rule(s) followed by applying the (sequence of) corresponding forward rule(s) yields the same result as applying the (sequence of) original TGG rule(s) assuming consistent matches [16,51].

Moreover, there is a correspondence between triple graphs belonging to the language of a given TGG and consistently and entirely marking transformation sequences via its forward rules. We formally state this correspondence as it is an ingredient for the proof of correctness of our synchronization algorithm.

**Lemma 1** (see [42, Fact 1] or [41, Lemma 4]) *Let a TGG  $GG$  be given. There exists a triple graph  $G = (G_S \leftarrow G_C \rightarrow G_T) \in \mathcal{L}(GG)$  if and only if there exists a transformation sequence like the one depicted in (1) via forward rules from  $GG$  such that  $G_0 = (G_S \leftarrow \emptyset \rightarrow \emptyset)$ ,  $G_t = (G_S \leftarrow G_C \rightarrow G_T)$ , and the transformation sequence is consistently and entirely marking.*

For practical purposes, forward rules and consistency patterns may be equipped with so-called *filter NACs* which can be automatically derived from the set of rules of the given TGG. The simplest examples of such filter NACs arise through the following analysis: For each rule that translate a node without translating adjacent edges it is first checked if other rules translate the same type of node but also translate an adjacent edge of some type. If this is the case, it is checked if there are further rules which only translate the detected kind of adjacent edge. If none is found, the original rule is equipped with a NAC forbidding the respective kind of edges. This avoids a dead-end in translation processes: In the presence of such a node with its adjacent edge, using the original rule to only translate the node leaves an untranslatable edge behind. The filter NAC of *Root-FWD-Rule* is derived in exactly this way. For the exact and more sophisticated derivation processes of filter NACs, we refer to the literature [29,35]. For our purposes it suffices to recall their distinguishing property: Filter NACs do not prevent “valid” transformation sequences of forward rules. We state this property in the terminology of our paper.

**Fact 1** ([29, Fact 4]) *Given a TGG  $GG = (\mathcal{R}, S)$ , for each  $r \in \mathcal{R}$ , let  $r^{FN}$  denote the corresponding forward rule that is additionally equipped with a set of derived filter NACs. (This set might be empty). For  $G_0 = (G_S \leftarrow \emptyset \rightarrow \emptyset)$ , there exists a consistently and entirely marking transformation sequence*

$$G_0 \Rightarrow_{r_1^F, m_1^F} G_1 \Rightarrow_{r_2^F, m_2^F} \dots \Rightarrow_{r_t^F, m_t^F} G_t$$

*via the forward rules (without filter NACs) derived from  $\mathcal{R}$  if and only if the sequence*

$$G_0 \Rightarrow_{r_1^{FN}, m_1^F} G_1 \Rightarrow_{r_2^{FN}, m_2^F} \dots \Rightarrow_{r_t^{FN}, m_t^F} G_t$$

*exists, i.e. if none of the filter NACs blocks one of the above rule applications.*

**Example 4** The source rules of the triple rules depicted in Fig. 2 are depicted in Fig. 4. They allow creating *Packages* and *Classes* on the source side without changing correspondence and target graphs. The formally existing empty graphs at correspondence and target sides are not depicted. The corresponding forward rules are given in Fig. 5. Their required elements are annotated with  $\boxtimes$  and their source marking elements with  $\square \rightarrow \boxtimes$ . The rule *Root-FWD-Rule* is equipped with a filter NAC: The given grammar does not allow creating a *Package* that is contained in another one with its original rule *Root-Rule*. Hence, the derived forward rule should not be used to translate a *Package*, which is contained in another one, to a *Folder*. As evident in the examples, the application of a source rule followed by the application of the corresponding forward rule amounts to the application of the original triple rule if matched consistently.

The consistency patterns that are derived from the TGG rules of our example are depicted in Fig. 9. They just check for existence of the pattern that occurs after applying the original TGG rule. A consistency pattern is equipped with the filter NACs of both its corresponding forward and backward rule. In our example, only *Root-Consistency-Pattern* receives such NACs; one from *Root-FWD-Rule* and the second one from the analogous backward rule. An occurrence of a consistency pattern in our example model indicates that a specific location corresponds to a concrete TGG rule application. Hence, a disappearance of such a match indicates that a former intact rule application has been broken and needs some fixing. We call this a *broken match for a consistency pattern* or, short, a *broken consistency match*. Practically, we will exploit an incremental pattern matcher to notify us about such disappearances.

### 3.3 Sequential independence

The proof of correctness of our synchronization approach relies on the notion of *sequential independence*. Transformations that are sequentially independent can be performed in arbitrary order.

**Definition 6** (*Sequential independence*) Given two transformation steps  $G \Rightarrow_{r_1, m_1} H_1 \Rightarrow_{r_2, m_2} X$ , via plain rules  $r_1, r_2$  these are *sequentially independent* if

$$n_1(R_1) \cap m_2(L_2) \subseteq n_1(K_1) \cap m_2(K_2) \tag{2}$$

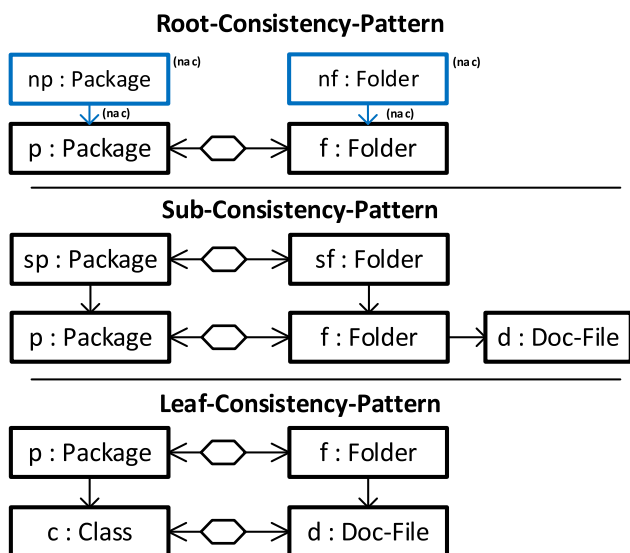


Fig. 9 Example: consistency patterns

where  $n_1$  is the comatch of the first transformation.

By the Local Church–Rosser Theorem [17, Theorem 3.20] the order of sequentially independent transformation can be switched. This means that, given a sequentially independent transformation sequence  $G \Rightarrow_{r_1, m_1} H_1 \Rightarrow_{r_2, m_2} X$ , there exists a sequentially independent transformation sequence  $G \Rightarrow_{r_2, m'_2} H_2 \Rightarrow_{r_1, m'_1} X$ . If  $r_1$  and  $r_2$  are equipped with NACs  $NAC_1$  and  $NAC_2$ , respectively, transformation steps as above are *sequentially independent* if condition (2) holds and moreover, the thereby induced matches  $m'_2 : L_2 \rightarrow G$  and  $m'_1 : L_1 \rightarrow H_2$  both satisfy the respective sets of NACs. In particular, the Local Church–Rosser Theorem still holds.

In our setting of graph transformation, it is easy to check the sequential independence of transformations [17,19]. A sequence  $t_1; t_2$  of two transformation steps is sequentially independent if and only if the following holds.

- $t_2$  does not match an element that  $t_1$  created.
- $t_2$  does not delete an element that  $t_1$  matches.
- $t_2$  does not create an element that  $t_1$  forbids.
- $t_1$  does not delete an element that  $t_2$  forbids.

## 4 Short-cut rules

Short-cut rules were introduced in [22] to take back an application of a TGG rule and to apply another one instead. This exchange of application shall be performed such that information loss is avoided. This means that model elements are checked for reuse before deleting them. We recall the construction of short-cut rules first and discuss their expressivity thereafter. Finally, we identify conditions for language-preserving applications of short-cut rules.

### 4.1 Construction of short-cut rules

We recall the construction of short-cut rules in a semiformal way and reuse an example of [22] for illustration; a formal treatment (in a category-theoretical setting) can be found in that paper. Given an inverse monotonic rule (i.e., a rule that purely deletes) and a monotonic rule, a short-cut rule combines their respective actions into a single rule. Its construction allows identifying elements that are deleted by the first rule as recreated by the second one. To motivate the construction, assume two monotonic rules  $r_1 : L_1 \rightarrow R_1$  and  $r_2 : L_2 \rightarrow R_2$  be given. Applying the inverse rule of  $r_1$  to a triple graph  $G$  provides an image of  $L_1$  in the resulting triple graph  $H$ . When applying  $r_2$  thereafter, the chosen match for  $L_2$  in  $H$  may intersect with the image of  $L_1$  yielding a triple graph  $L_\cap$ . This intersection can also be understood as saying that  $L_\cap$  provides a partial match for  $L_2$ . The inverse application of the first rule deletes elements which may be recreated again. In this case, it is possible to extend the subtriple graph  $L_\cap$  of  $H$  to a subtriple graph  $R_\cap$  of  $H$  with these elements. In particular,  $R_\cap$  is a subtriple graph of  $R_1$  and  $R_2$  as it includes elements only that have been deleted by the first rule and created by the second. Based on this observation, the construction of short-cut rules is defined as follows (slightly simplified and directly merged with an example):

**Construction 7 (Short-cut rule)** Let two plain, monotonic rules  $r_1 = L_1 \rightarrow R_1$  and  $r_2 = L_2 \rightarrow R_2$  be given. A short-cut rule  $r_{sc}$  for the rule pair  $(r_1, r_2)$ , where  $r_1$  is considered to be applied inversely, is constructed in the following way:

1. *Choice of common kernel:* A (potentially empty) subtriple graph  $L_\cap$  of  $L_1$  and  $L_2$  and a subtriple graph  $R_\cap$  of  $R_1$  and  $R_2$  with  $L_\cap \subseteq R_\cap$  are chosen. We call  $L_\cap \subseteq R_\cap$  a *common kernel* of both rules. In Fig. 10, an example of such a common kernel is given. It is a common kernel for rule pair *(Root-Rule, Sub-Rule)*. The common kernel is depicted in the center of Fig. 10. This choice of a common kernel will lead to *Connect-Root-SC-Rule* as resulting short-cut rule. In this example,  $L_\cap$  is empty and  $R_\cap$  extends  $L_\cap$  by identifying the *Packages* `p`, *Folder* `f`, and the correspondence node in between. The elements of  $R_\cap \setminus L_\cap$ , called *recovered elements*, are to become the elements that are preserved by an application of the short-cut rule compared to reversely applying the first rule followed by applying the second one (provided that these applications overlap in  $L_\cap$ ). In the example case, the whole graph  $R_\cap$  is recovered as  $L_\cap$  is empty.
2. *Construction of LHS and RHS:* One first computes the union  $L_\cup$  of  $L_1$  and  $L_2$  along  $L_\cap$ . The result is then united with  $R_1$  along  $L_1$  and  $R_2$  along  $L_2$ , respectively,





### 4.2 Expressivity of short-cut rules

Given a set of rules, there are *two degrees of freedom* when deciding which short-cut rules to derive from them: First, one has to choose for which *pairs of rules* short-cut rules shall be derived. Secondly, given a pair of rules, there is typically not only one way to construct a short-cut rule for them: In general, there are different choices for a *common kernel*. However, when fixing a common kernel, i.e.,  $L_{\cap}$  and  $R_{\cap}$ , the result of the construction is uniquely determined. If, moreover, the LHSs and RHSs of the rules are finite, the set of possible common kernels is finite as well.

As short-cut rules correspond to possible (complex) edits of a triple graph, the more short-cut rules are derived, the more user edits are available which can directly be propagated by the corresponding repair rules. But the number of rules that has to be computed (and maintained throughout the synchronization process) in this way, would quickly grow. And maybe several of the constructed rules would capture edits that are possible in principle but unlikely to ever be performed in a realistic scenario. Hence, some trade-off between expressivity and maintainability has to be found.

We shortly discuss these effects of choices: The construction of short-cut rules is defined for *any two* monotonic rules [22]—we do not need to restrict to the rules of a given TGG but may also use monotonic rules that have been constructed as so-called *concurrent rules* [17] of given TGG rules as input for the short-cut rule construction. A concurrent rule combines the actions of two (or more) subsequent rule applications into a single rule. Hence, deriving short-cut rules from concurrent rules that have been built of given TGG rules leads to short-cut rules that capture even more complex edits into a single rule. The next example presents such a derived short-cut rule. While our conceptual approach is easily extended to support such rules, we currently stick with short-cut rules directly derived from a pair of rules of the given TGG in our implementation.

**Example 6** The short-cut rule *Delete-Middle-SC-Rule* depicted in Fig. 13 is not directly derived of the TGG rules depicted in Fig. 2. Instead, the concurrent rule of two given applications of *Sub-Rule* is constructed first. This concurrent rule directly creates a chain of two *Packages* and *Folder* into an existing pair of *Package* and *Folder*. The rule in Fig. 13 is a short-cut rule of this concurrent rule and *Sub-Rule*. It takes back the creation of a chain such that the bottom package is directly included in the top package in Fig. 13.

Concerning the choice of a common kernel, we follow two strategies. In both strategies, we overlap as many of the newly created elements of the two input rules as possible since these are the elements that we try to preserve.

A *minimal overlap* overlaps created elements only, i.e., no context elements. An example is *Sub-Rule*, which overlapped

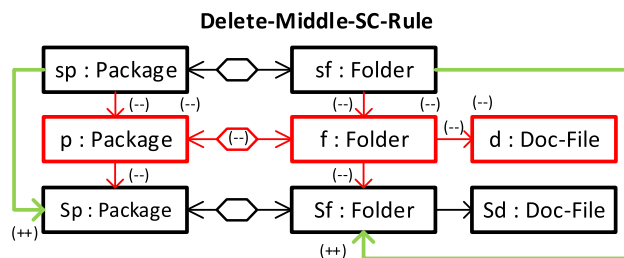


Fig. 13 Example for a short-cut rule not directly derived from the rules of our example TGG

with itself, results in *Move-To-New-Sub-SC-Rule* and which corresponds to a move refactoring step.

A *maximal overlap* overlaps not only created elements of both rules but also context elements. Creating such an overlap for *Sub-Rule* with itself would result in the *Sub-Consistency-Pattern*, which has no effect when applied. However, when overlapping different rules with each other, it is often useful to re-use context elements. This is the case, for example, for *VariableDec-2-Parameter-Rule* and *TypeAccess-2-ReturnType-Rule* of our evaluation rule set in Fig. 19 in the “Appendix” below. A full overlap between both rules would allow to transform a signature parameter to a return parameter of the same method and of the same type and, vice versa.

Both strategies aim to create different kinds of short-cut rules with specific purposes. Since generating all possible overlaps and thus short-cut rules is expensive, we chose a heuristic approach to generate a useful subset of them.

As we are dealing with triple graphs being composed of source, target and correspondence graphs, the overlap of source graphs should correspond to that of target graphs. This restricts the kind of “reuse” of elements the derived short-cut rules enable. The allowance of any kind of overlap may include unintended ones. We argue for the usefulness of these strategies in our evaluation in Sect. 8.

### 4.3 Language preserving short-cut rule applications

The central intuition behind the construction of short-cut rules is to replace the application of a monotonic triple rule by another one. In this sense, a short-cut rule captures a complex edit operation on triple graphs that (in general) cannot be performed directly using the rules of a TGG. We illustrate this behavior in the following. Subsequently, we discuss the circumstances under which applications of short-cut rules are “legal” in the sense that the result still belongs to the language of the respective TGG.

Let a TGG  $GG$  and a sequence of transformations

$$G_0 \Rightarrow_{r_1, m_1} G_1 \Rightarrow_{r_2, m_2} G_2 \Rightarrow \dots \Rightarrow_{r_t, m_t} G_t \tag{3}$$

be given where all the  $r_i$ ,  $1 \leq i \leq t$ , are rules of  $GG$ , all the  $m_i$  denote the respective matches, and  $G_0 \in \mathcal{L}(GG)$ ; in particular  $G_t \in \mathcal{L}(GG)$  as well. Fixing some  $j \in \{1, \dots, t\}$  and some rule  $r$  of  $GG$ , we construct a short-cut rule  $r_{sc}$  for  $(r_j, r)$  with some common kernel  $L_\cap \subseteq R_\cap$ . Next, we can consider the transformation sequence

$$G_0 \Rightarrow_{r_1, m_1} G_1 \Rightarrow_{r_2, m_2} G_2 \Rightarrow \dots \Rightarrow_{r_t, m_t} G_t \Rightarrow_{r_{sc}, m_{sc}} G'_t$$

that arises by appending an application of  $r_{sc}$  to transformation sequence (3). Under certain technical circumstances (which we will state below) this transformation sequence is equivalent<sup>6</sup> to the sequence

$$\begin{aligned} G_0 \Rightarrow_{r_1, m_1} G_1 \Rightarrow \dots \Rightarrow_{r_{j-1}, m_{j-1}} G_{j-1} \Rightarrow_{r, m'_{sc}} G'_j \\ \Rightarrow_{r_{j+1}, m'_{j+1}} \dots \Rightarrow_{r_t, m'_t} G''_t \end{aligned} \quad (4)$$

where the application of  $r_j$  at match  $m_j$  is replaced by an application of  $r$  at a match  $m'_{sc}$  that is derived from the match  $m_{sc}$  of the short-cut rule. The following matches  $m_{j+1}, \dots, m_t$  have been adapted accordingly. They still match the same elements but formally they do so in other triple graphs. In particular,  $G'_t$ , the result of the transformation sequence (4), is isomorphic to  $G'_j$  and hence,  $G'_t$  can be understood as arising by replacing the  $j$ -th rule application in the transformation sequence (3) by an application of the rule  $r$ ; thus,  $G'_t$  also belongs to the language of the TGG: The sequence (4) starts at a triple graph  $G_0 \in \mathcal{L}(GG)$  and solely consists of applications of rules from  $GG$ .

**Example 7** Consider the triple graph depicted in Fig. 3(a). It arises by applying *Root-Rule*, followed by two applications of *Sub-Rule*, and finally an application of *Leaf-Rule*. When matched as already described in the introductory example, an additional application of *Make-Root-SC-Rule* to this triple graph results in the one depicted in Fig. 3(c). Alternatively, this can be derived by two applications of *Root-Rule*, followed by an application of *Sub-Rule* and *Leaf-Rule* each. As schematically depicted in Fig. 14, the application of the short-cut rule *Make-Root-SC-Rule* transforms a transformation sequence deriving the first triple graph into a transformation sequence deriving the second one by replacing an application of *Sub-Rule* by one of *Root-Rule*.

In the following, we state when the above described behavior is the case (in a somewhat less technical language than originally used).

<sup>6</sup> The formal notion of equivalence used here is called *switch equivalence* and captures the idea that, in case of sequential independence, the order of rule applications might be switched while using basically the same match for each rule application and receiving the same result; compare, e.g., [8,38].

**Theorem 2** ([23, Theorem 8]) *Let the transformation sequencex (3) be given and let  $r_{sc}$  be a short-cut rule that is derived from  $(r_j, r)$ . If the following three conditions are met, this sequence is equivalent to sequence (4) where original TGG rules are applied only.*

1. Reversing match: *The application of  $r_{sc}$  at  $m_{sc}$  reverses the application of  $r_j$ , i.e.,  $n_j(R_j) = m_{sc}|_{R_j}(R_j)$ .*
2. Sequential independence:
  - (a) Non-disabling match: *The application of  $r_{sc}$  at  $m'_{sc}$  does not delete elements used in the applications of  $r_{j+1}, \dots, r_t$ .*
  - (b) Context-preserving match: *The match  $m_{sc}$  for  $r_{sc}$  already exists in  $G_{j-1}$ . Since the assumption on the match to be reversing already ensures this for elements of  $L_{sc}$  that stem from  $R_j$ , context-preservation ensures in particular that the presumed elements of  $r_{sc}$  are matched to elements already existing in  $G_{j-1}$ .*


**Example 8** We illustrate each of the above mentioned conditions:

1. *Reversing match*: In our example of matching *Connect-Root-SC-Rule* to the triple graph (c) in Fig. 3 this means that its nodes  $\mathfrak{p}$  and  $\mathfrak{f}$  (and the correspondence node in between) are allowed to be matched to elements only that have been created using *Root-Rule*. In this way, it is avoided to misuse the rule to introduce *Packages* (and *Folder*) that are contained by more than one *Package* (or *Folder*).
2. *Non-disabling match*: For example, *Delete-Middle-SC-Rule* from Fig. 13 is not allowed to delete *Packages* and *Folder* that already contain *Classes* or *Doc-Files*, respectively.
3. *Context preserving match*: Returning to our example of matching *Connect-Root-SC-Rule* to the triple graph (c) in Fig. 3 this means that as soon as nodes  $\text{subP}$  and  $\text{subF}$  in that triple graph have been chosen as matches for the nodes  $\mathfrak{p}$  and  $\mathfrak{f}$  of *Connect-Root-SC-Rule*, the nodes  $\text{leafP}$  and  $\text{leafF}$  are not allowed to be chosen as matches for nodes  $\mathfrak{sp}$  and  $\mathfrak{sf}$  of *Connect-Root-SC-Rule*. The creation of  $\text{leafP}$  and  $\text{leafF}$  depends on  $\text{subP}$  and  $\text{subF}$  being created first. In this way, the introduction of cyclic dependencies between elements is avoided.

## 5 Constructing language-preserving repair rules

In this section, we formally define the derivation of repair rules from a given TGG and characterize valid applications of these. Our general idea is to construct *repair rules* that

**Fig. 14** Example: transforming sequences of rule applications by applying short-cut rules



can be used during model synchronization processes that are based on the formalism of TGGs. Our construction of such repair rules is based on *short-cut rules* which we recalled in Sect. 4.

### 5.1 Deriving repair rules from short-cut rules

Having defined short-cut rules, they can be operationalized to get edit rules for source graphs and forward rules that repair these edits. As such edits may delete source elements, correspondence elements may be left without corresponding source elements. Hence, the resulting triple graphs show a form of partiality. They are called *partial triple graphs*. Given a model, formally considered as triple graph  $G_S \xleftarrow{\sigma_G} G_C \xrightarrow{\tau_G} G_T$ , a user edit on  $G_S$  may consist of the deletion and/or creation of graph elements, resulting in a graph  $G'_S$ . In general, the “old” correspondence morphism  $\sigma_G : G_C \rightarrow G_S$  does not extend to a correspondence morphism from  $G_C$  to  $G'_S$ : The user might have deleted elements in the image of  $\sigma_G$ . However, there is a *partial morphism*  $\sigma'_G : G_C \dashrightarrow G'_S$  that is defined for all elements whose image under  $\sigma_G$  still exists.

**Definition 8** (*Partial triple graph*) A *partial graph morphism*  $f : A \dashrightarrow B$  is a graph morphism  $f : A' \rightarrow B$  where  $A'$  is a subgraph of  $A$ ;  $A'$  is called the *domain* of  $f$ .

A *partial triple graph*  $G' = G'_S \xleftarrow{\sigma'_G} G'_C \xrightarrow{\tau'_G} G'_T$  consists of three graphs  $G'_S, G'_C, G'_T$  and two partial graph morphisms  $\sigma'_G : G'_C \dashrightarrow G'_S$  and  $\tau'_G : G'_C \dashrightarrow G'_T$ .

Given a triple graph  $G = (G_S \xleftarrow{\sigma_G} G_C \xrightarrow{\tau_G} G_T)$  and a user edit of  $G_S$  that results in a graph  $G'_S$ , the *partial triple graph induced by the edit* is  $G'_S \xleftarrow{\sigma'_G} G_C \xrightarrow{\tau_G} G_T$  where  $\sigma'_G$  is obtained by restricting  $\sigma_G$  to those elements  $x$  of  $G_C$  (node or edge) for which  $\sigma_G(x) \in G_S$  is still an element of  $G'_S$ .

According to the above definition, triple graphs are special partial triple graphs, namely those, where the domain of both partial correspondence morphisms is the whole correspondence graph  $G_C$ .

When operationalizing short-cut rules, i.e., splitting them into a *source* and a *forward rule*, we also have to deal with this kind of partiality: In contrast to the rules of a given TGG, a short-cut rule might delete an element. Hence, its forward rule might need to contain a correspondence element for

which the corresponding source element is missing; it is referenced in the short-cut rule. This element is deleted by the corresponding source rule.

**Definition 9** (*Source and forward rule of short-cut rule. Repair rule*) Given a pair  $(r_1, r_2)$  of plain, monotonic triple rules with short-cut rule  $r_{sc} = (L_{sc} \xleftarrow{l_{sc}} K_{sc} \xrightarrow{r_{sc}} R_{sc})$ , the *source* and *forward rule* of  $r_{sc}$  are defined as

$$r_{sc}^S := (L_{sc}^S \xleftarrow{(l_{sc,S}, id_\emptyset, id_\emptyset)} K_{sc}^S \xrightarrow{(r_{sc,S}, id_\emptyset, id_\emptyset)} R_{sc}^S)$$

and

$$r_{sc}^F := (L_{sc}^F \xleftarrow{(id_{R_{sc,S}}, l_{sc,C}, l_{sc,T})} K_{sc}^F \xrightarrow{(id_{R_{sc,S}}, r_{sc,C}, r_{sc,T})} R_{sc}^F)$$

where

$$\begin{aligned} L_{sc}^S &:= (L_{sc,S} \leftarrow \emptyset \rightarrow \emptyset), \\ K_{sc}^S &:= (K_{sc,S} \leftarrow \emptyset \rightarrow \emptyset), \\ R_{sc}^S &:= (R_{sc,S} \leftarrow \emptyset \rightarrow \emptyset), \\ L_{sc}^F &:= (R_{sc,S} \leftarrow L_{sc,C} \rightarrow L_{sc,T}), \\ K_{sc}^F &:= (R_{sc,S} \leftarrow K_{sc,C} \rightarrow K_{sc,T}), \text{ and} \\ R_{sc}^F &:= (R_{sc,S} \leftarrow R_{sc,C} \rightarrow R_{sc,T}). \end{aligned}$$

Given a TGG  $GG$ , a *repair rule* for  $GG$  is the forward rule  $r_{sc}^F$  of a short-cut rule  $r_{sc}$  where  $r_{sc}$  has been constructed from a pair of rules of  $GG$ .

For more details (in particular, the definition of morphisms between partial triple graphs), we refer the interested reader to the literature [23,37]. In this paper, we are more interested in conveying the intuition behind these rules by presenting examples. We next recall the most important property of this operationalization, namely that, as in the monotonic case, an application of a short-cut rule corresponds to the application of its source rule, followed by an application of the forward rule if consistently matched.

**Theorem 3** ([23, Theorem 7] and [37, Theorem 23]) *Given a short-cut rule  $r_{sc}$ , there is a transformation*

$$(G_S \leftarrow G_C \rightarrow G_T) \Rightarrow_{r_{sc}, m_{sc}} (H_S \leftarrow H_C \rightarrow H_T)$$





- (a) *Non-disabling match*: The application of  $r_{sc}^F$  does not delete elements used in the applications of  $r_{j+1}^{FN}, \dots, r_t^{FN}$  nor does it create elements forbidden by one of the filter NACs of those forward rules.
- (b) *Context-preserving match*: The presumed source elements of the repair rule  $r_{sc}^F$  (which accord to the presumed source elements of the short-cut rule  $r_{sc}$ ) are matched to elements of  $H_S$  which are marked as translated in  $G_{j-1,S}$ . Presumed context and target elements of  $r_{sc}^F$  are matched to elements of  $G_{t,C}$  and  $G_{t,T}$  that are already created in  $G_{j-1,C}$  and  $G_{j-1,T}$ , respectively. This means, elements stemming from the LHS  $L$  of  $r$  which have not been identified with elements from  $L_j$  in the short-cut rule  $r_{sc}$  are matched to elements already translated/existing in  $G_{j-1}$ .

Together, items (a) and (b) imply that the application of  $r_{sc}^F$  is sequentially independent from each of the applications of  $r_k^{FN}$  for  $j + 1 \leq k \leq t$ .

3. *Creation-preserving match*: All source elements that are newly created by short-cut rule  $r_{sc}$ , i.e., the source elements of  $R_S \setminus L_S$  that have not been merged with an element of  $R_{j,S} \setminus L_{j,S}$  during the short-cut rule construction, are matched to elements which are yet untranslated in  $G_{t,S}$ .

The following corollary uses Theorem 3 to transfer the statement of Theorem 2 to repair rules. The additional requirement on the match to be creation preserving in the above definition of valid matches (compared to Theorem 2 for short-cut rules) originates from the fact that forward rules do not create but mark source elements.

**Corollary 1** *Let a TGG  $GG$  and a consistently marking transformation sequence as in (5), followed by an edit step exactly as in Definition 10 above be given. Then, applying  $r_{sc}^F$  at a valid match  $m_{sc}^F$  in  $G'$  induces a consistently marking transformation sequence*

$$G'_0 \Rightarrow_{r_1^{FN}, m_1^F} G'_1 \Rightarrow_{r_2^{FN}, m_2^F} \dots \Rightarrow_{r_{j-1}^{FN}, m_{j-1}^F} G'_{j-1} \Rightarrow_{r_{sc}^F, m_{sc}^F} X \tag{6}$$

with  $G'_i = (H_S \leftarrow G_{i,C} \rightarrow G_{i,T})$  for  $0 \leq i \leq j - 1$ .

**Proof** For a valid match  $m_{sc}^F$  of  $r_{sc}^F$ , by its reversing property, the conditions of Theorem 3 are met. Hence, we obtain a sequence

$$G_0 \Rightarrow_{r_1^{FN}, m_1^F} \dots \Rightarrow_{r_t^{FN}, m_t^F} G_t \Rightarrow_{r_{sc}, m_{sc}} X'$$

As a consistently marking sequence of forward rules corresponds to a sequence of TGG rule applications, and the preconditions of Theorem 2 are met (“exists” is exchanged by

“marked” on the source component), this sequence induces a sequence

$$G_0 \Rightarrow_{r_1^{FN}, m_1^F} \dots \Rightarrow_{r_{j-1}^{FN}, m_{j-1}^F} G_{j-1} \Rightarrow_{r, m} X$$

(where we do not care for the further applications of forward rules).

Now, we can split  $r$  into its source and forward rule. Its source rule is sequentially independent from the other forward rule applications:  $r_{sc}^S$  does not delete anything, the rules  $r_1^{FN}, \dots, r_{j-1}^{FN}$  match, and does not create a filter NAC violation by assumption and, as a consequence,  $r^S$  does not. Hence, by the local Church–Rosser Theorem, we might equivalently switch the application of  $r^S$  to the beginning of the sequence and obtain sequence (6), as desired. Moreover, by Lemma 2, the filter NAC of  $r^F$  holds whenever  $m_{sc}^F$  satisfies the filter NAC of  $r_{sc}^F$ .

Finally, as the start of the transformation sequence (up to index  $j - 1$ ) is context preserving, and by assumption 2. (b), the match  $m_{sc}^F$  matches presumed elements of  $r_{sc}^F$  to already translated ones (in  $H_S$ ) or already created ones (in  $G_{j-1,C}$  and  $G_{j-1,T}$ ), this sequence is context preserving. Analogously, assumption 3. ensures that it is creation-preserving: No element which is already marked as translated in  $G_{t,S}$  is marked a second time. Hence, the whole sequence is consistently marking.  $\square$

## 6 Synchronization algorithm

In this section, we discuss our synchronization algorithm that is based on the correct application of derived repair rules. We first present the algorithm and consider its formal properties subsequently. The section closes with a short example for a synchronization based on our algorithm and a discussion of extensions and support for advanced TGG features.

### 6.1 The basic setup

We assume a TGG  $GG$  with plain, monotonic rules to be given. Its language defines consistency. This means that a triple graph  $G = (G_S \leftarrow G_C \rightarrow G_T)$  is consistent if and only if  $G \in \mathcal{L}(GG)$ .

The problem A consistent triple graph  $G = (G_S \leftarrow G_C \rightarrow G_T) \in \mathcal{L}(GG)$  is given; by Lemma 1 there exists a corresponding consistently and entirely marking sequence  $t$  of forward rule applications. After editing source graph  $G_S$  we get  $G' = (H_S \leftarrow G_C \rightarrow G_T)$ . Generally, the result  $G'$  is a partial triple graph and does not belong to  $\mathcal{L}(GG)$ . We assume that all the edits are performed by applying source rules. They may be derived from the original TGG rules or from short-cut rules. Our goal is to provide a model synchro-

nization algorithm that, given  $G = (G_S \leftarrow G_C \rightarrow G_T) \in \mathcal{L}(GG)$  and  $G' = (H_S \leftarrow G_C \rightarrow G_T)$  as input, computes a triple graph  $H = (H_S \leftarrow H_C \rightarrow H_T) \in \mathcal{L}(GG)$ . As a side condition, we want to minimize the amount of elements of  $G_C$  and  $G_T$  that are deleted and recreated during that synchronization.

*Ingredients of our algorithm* We provide a rule-based model synchronization algorithm leveraging an incremental pattern matcher. During that algorithm, rules are applied to compute a triple graph  $(H_S \leftarrow H_C \rightarrow H_T) \in \mathcal{L}(GG)$  from the (partial) triple graph  $(H_S \leftarrow G_C \rightarrow G_T)$ . We apply two different kinds of rules, namely

1. forward rules derived from the rules of the TGG  $GG$  and
2. repair rules, i.e., operationalized short-cut rules.

Forward rules serve to propagate the addition of elements. The use of these rules for model synchronization is standard. However, the use of additional repair rules and the way in which they are employed are conceptually novel.<sup>7</sup> The repair rules allow directly propagating more complex user edits.

During the synchronization process, the rules are applied reacting to notifications by an incremental pattern matcher. We require this pattern matcher to provide the following information:

1. The original triple graph  $G = (G_S \leftarrow G_C \rightarrow G_T)$  is covered with consistency patterns. When considering the induced matches for forward rules, every element of  $G_S$  is marked exactly once. The dependency relation between elements required by these matches is acyclic. This means that the induced transformation sequence of forward rules is consistently and entirely marking. Such a sequence always exists since  $G \in \mathcal{L}(GG)$ ; see Lemma 1.
2. Broken consistency matches are reported. A match for a consistency pattern in  $G$  is broken in  $G'$  if one of the elements it matches or creates has been deleted or if an element has been created that violates one of the filter NACs of that consistency pattern.
3. The incremental pattern matcher notifies about newly occurring matches for forward rules. It does so in a correct way, i.e., it only notifies about matches that lead to consistently marking transformations.
4. In addition, the incremental pattern matcher informs a precedence graph. This precedence graph contains information about the particular sequential dependencies of the elements in the partial triple graph. Here, an element is dependent on another one if the forward rule application marking the former matches the latter element as

required. We consider the transitive closure of this relation.

## 6.2 Synchronization process

Our synchronization process is depicted in Algorithm 1. It applies rules to translate elements and repair rule applications. In that, it applies a different strategy than suggested in [41,42]. There, invalid rule applications are revoked as long as there exist any. Subsequently, forward rules are applied as long as possible. By trying to apply a suitable repair rule instead of revoking an invalid rule application, we are able to avoid deletion and recreation of elements. Our synchronization algorithm is defined as follows. Note that we present an algorithm for synchronizing in forward direction (from source to target) while synchronizing backwards is performed analogously.

The function *synchronize* is called on the current partial triple graph that is to be synchronized. In line 2, *updateMatches* is called on this partial triple graph. It returns the set of consistency matches currently broken, a set of consistency matches being still intact, and a set of forward TGG rule matches.

By calling the function *isFinished* (line 4), termination criteria for the synchronization algorithm are checked. If the set of broken consistency matches and the set of forward TGG rule matches are both empty and all elements of the source graph are marked as translated, the synchronization algorithm terminates (line 18). Yet, if both sets are empty but there are still untranslated elements in the source graph, an exception is thrown in line 20, signaling that the (partial) triple graph is in an inconsistent state.

Subsequently, function *translate* is used (line 7) to propagate the creation of elements: If the set of forward TGG rule matches is non-empty (line 24), we choose one of these matches, apply the corresponding rule, and continue the synchronization process (line 27). This step is done prior to any repair. The purpose is to create the context which may be needed to make repair rules applicable. An example for such a context creation is the insertion of a new root *Package* which has to be translated into a root *Folder* before applying *Connect-Root-Repair-Rule* thereafter (see Fig. 5). If the above cases do not apply, there must be at least one broken consistency match and the corresponding rule application has to be repaired (line 10): Hence, we choose one broken consistency match (line 32) for which a set of suitable repair rules is determined. A broken consistency match includes information about the rule it corresponds to (e.g., the name of the rule). Furthermore, it includes which elements are missing or which filter NACs are violated such that the corresponding application does not exist any more. We calculate the set of matches of *repair rules* (i.e., forward short-cut rules) that stem from short-cut rules revoking exactly the rule that

<sup>7</sup> Note that consistency is still defined by the (plain, monotonic) rules of the given TGG; the general repair rules are derived only to improve the synchronization process.

**Algorithm 1** eMoflon – Synchronization Process

---

```

1: function SYNCHRONIZE(tripleGraph)
2:   (brokenCMatches, intactCMatches, fwdMatches) ← updateMatches(tripleGraph)
3:
4:   if isFinished(tripleGraph, fwdMatches, brokenCMatches) then
5:     return
6:
7:   if translate(tripleGraph, fwdMatches) then
8:     return
9:
10:  (cMatch, success) ← repair(tripleGraph, brokenCMatches)
11:  if !success then
12:    throw InconsistentStateException
13:  return
14:
15:  function ISFINISHED(tripleGraph, fwdMatches, brokenCMatches)
16:  if isEmpty(brokenCMatches) && isEmpty(fwdMatches) then
17:    if allElementsTranslated(tripleGraph.source) then
18:      return true
19:    else
20:      throw InconsistentStateException
21:  return false
22:
23:  function TRANSLATE(tripleGraph, fwdMatches)
24:  if !isEmpty(fwdMatches) then
25:    fwdMatch = chooseMatch(fwdMatches)
26:    tripleGraph ← applyRule(tripleGraph, fwdMatch, getFWDRule(fwdMatch))
27:    synchronize(tripleGraph)
28:    return true
29:  return false
30:
31:  function REPAIR(tripleGraph, brokenCMatches)
32:  cMatch ← chooseMatch(brokenCMatches)
33:  scRules ← getSuitableSCRules(cMatch)
34:  scMatches ← findSCMatches(scRules, cMatch)
35:
36:  while !isEmpty(scMatches) do
37:    scMatch ← chooseMatch(scMatches)
38:    if isValidMatch(scMatch) then
39:      tripleGraph ← applyRule(tripleGraph, cMatch, getSCRule(scMatch))
40:      synchronize(tripleGraph)
41:      return (cMatch, true)
42:  return (cMatch, false)

```

---

corresponds to the broken consistency match. In particular, by knowing which elements of a broken rule application still exist in the current source graph, we can stick to those repair rules that preserve exactly the still existing elements.

While the calculated set of unprocessed *repair rule* matches is not empty (line 36), we choose one of these matches and check whether it is *valid*. By constructing the partial match of a *repair rule*, we only need to ensure that none of its *presumed elements* is matched in such a way that a cyclic dependency is introduced. This means that they must not be matched to elements that are dependent of elements to which the *recovered elements* are matched. If a match is valid, we apply the corresponding repair rule and continue the synchronization process (line 40). If no such rule or valid match is available, an exception is thrown (line 12).

### 6.3 Formal properties of the synchronization process

We discuss the termination, correctness, and completeness of our synchronization algorithm.

Our algorithm terminates as long as every forward rule translates at least one element (which is a quite common condition; compare [30, Lemma 6.7] or [41, Theorem 3]).

**Theorem 4** *Let a TGG  $GG$  with plain, monotonic rules be given. If every derived forward rule of  $GG$  has at least one source marking element, our algorithm terminates for any finite input  $G' = (H'_S \leftarrow G_C \rightarrow G_T)$ .*

**Proof** The algorithm terminates—by either throwing an exception or returning a result—if at one point both, the set of broken consistency matches and the set of matches for

forward rules are empty; compare the function *isFinished* starting in line 15.

The algorithm is called recursively, always applying a forward rule if a match is available. As every forward rule marks at least one element as translated and forward rules are only matched in such a way that source marking elements are matched to yet untranslated ones, the application of forward rules (lines 24 et seq.), i.e., the recursive call of function *translate*, halts after finitely many steps. Moreover, an application of a forward rule never introduces a new broken consistency match: As it neither creates nor deletes elements in the source graph, it cannot delete elements matched by a consistency pattern nor create elements forbidden by one. This means that, as soon as the set of broken consistency matches is empty, the whole synchronization algorithm will terminate. We show that at some point this set of broken consistency matches will be empty or an exception is thrown.

Whenever the algorithm is called with an empty set of matches for forward rules, broken consistency matches are considered by applying a repair rule, i.e., by calling the function *repair*. New matches for forward rules can result from this; as discussed above, newly appearing matches for forward rules are unproblematic. However, an application of a repair rule does not introduce a new violation of any consistency match: As it does not create source elements, it cannot introduce violations of filter NACs. And by the condition on valid matches to be non-disabling (condition 2. (a) in Definition 10), no elements needed by other consistency matches are deleted. Hence, by application of a repair rule, the number of invalid consistency matches is reduced by one and the algorithm terminates as soon as all broken consistency matches are repaired. If there is a broken consistency match that cannot be repaired—either because no suitable repair rule or no valid match is available—an exception is thrown and the algorithm stops.  $\square$

*Correctness* Upon termination without exception, our algorithm is correct.

**Theorem 5** (Correctness of algorithm) *Let a TGG  $GG$  with plain, monotonic rules, a triple graph  $G = (G_S \leftarrow G_C \rightarrow G_T) \in \mathcal{L}(GG)$ , and a partial triple graph  $G' = (G'_S \leftarrow G_C \rightarrow G_T)$  that arises by a user edit step on the source graph be given. If our synchronization algorithm terminates without exception and yields  $H = (H_S \leftarrow H_C \rightarrow H_T)$  as output, then  $H_S = G'_S$  and  $H \in \mathcal{L}(GG)$ .*

**Proof** We see immediately that  $H_S = G'_S$  since none of the applied rules modifies the source graph. If the synchronization process terminates without exception, all elements are translated, no matches for forward rules are found, and no consistency match is broken any more. This means that

the collected matches of the forward rules form an entirely marking transformation sequence. By Lemma 1, we have to show that this sequence is also consistently marking. Then, the matches of the forward rules that correspond to the matches of the consistency patterns that the incremental pattern matcher has collected encode a transformation sequence that allows translating the triple graph  $(H_S \leftarrow \emptyset \rightarrow \emptyset)$  to a triple graph  $(H_S \leftarrow H_C \rightarrow H_T) \in \mathcal{L}(GG)$ . We assume that the incremental pattern matcher recognizes all broken consistency matches and reports correct matches for forward rules only. This means, throughout the application of forward rules, the set of all valid consistency matches remains consistently marking. We have to show that this is also the case for repair rule applications. If it is, upon termination without exception, there is an entirely and consistently marking sequence of forward rules which corresponds to a triple graph from  $GG$  by Lemma 1.

Whenever we apply a repair rule we are (at least locally) in the situation of Corollary 1: There is a (maybe empty) sequence of consistently marking forward rule applications and a suitable broken consistency pattern indicates, that a user edit step applying the source rule  $r_{sc}^S$  of a short-cut rule  $r_{sc}$  has taken place. Applying the repair rule  $r_{sc}^F$  at a valid match amounts to replacing the application of rule  $r_j^F$ , whose consistency pattern was broken, by rule  $r^F$  in a consistently marking way.  $\square$

We only informally discuss *completeness*. We understand completeness as follows: for every input  $G' = (H_S \leftarrow G_C \rightarrow G_T)$  with  $H_S \in \mathcal{L}_S(GG)$ , we obtain a result  $H = (H_S \leftarrow H_C \rightarrow H_T) \in \mathcal{L}(GG)$ . In general, the above proposed algorithm is not complete. We randomly apply forward rules at available matches (without using backtracking) but the choice and order of such applications can affect the result if the final sequence of forward rule applications leads to a dead-end or translates the given source graph. However, the algorithm is complete whenever the set of forward rules is of such a form that the order of their application does not make a difference (somewhat more formally: they meet some kind of confluence) and the user edit is of the form discussed in Sect. 6.1. Analogous restrictions on forward rules hold for other synchronization processes that have been formally examined for completeness [30,41]. Adding filter NACs to the forward rules of a TGG is a technique that can result in such a set of confluent forward rules even if the original set of forward rules is not. Moreover, there are static methods to test TGGs for such a behavior [6,30]; they check for sufficient but not for necessary criteria. If it is known that the set of forward rules of a given TGG guarantees completeness and the edit is of a suitable kind, a thrown exception during our synchronization process implies that  $H_S \notin \mathcal{L}_S(GG)$ .



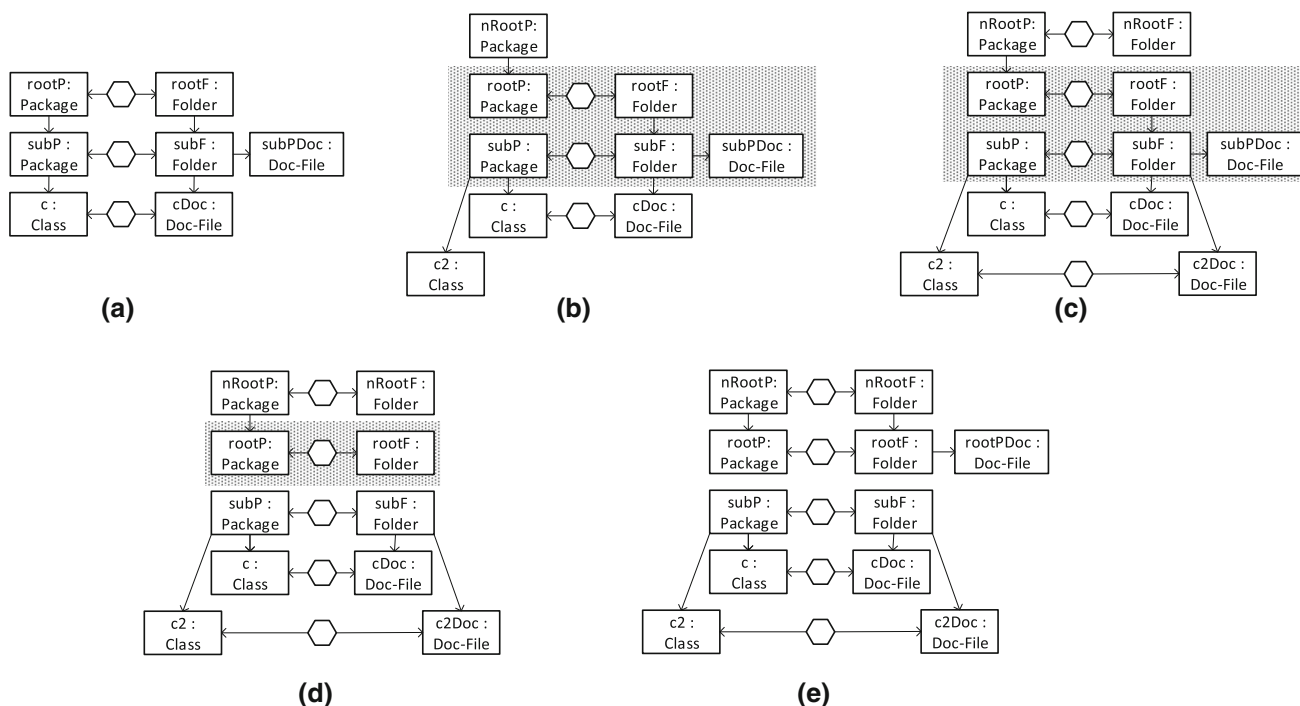


Fig. 16 Example of our proposed synchronization algorithm. Grey background indicates broken consistency matches

### 6.4 A synchronization example

We illustrate our synchronization algorithm with an example illustrated in Fig. 16. For simplicity, we neglect the content attribute and concentrate on the structural behavior. As a starting point, we assume that a user edits the source graph of the triple graph depicted in Fig. 16(a) (in the following, we will refer to the triple graphs occurring throughout the algorithm just by their numbers). She adds a new root package above `rootP`, removes the link between `Packages` `rootP` and `subP`, and creates a further class `c2`. All these changes are specified by either a source rule of the TGG or the source rule of a derived short-cut rule. The resulting triple graph is depicted in (b). The elements in front of the grey background are considered to be inconsistent, due to a broken consistency match. Furthermore, `c2` and `nRootP` are not translated, yet. In the first two passes of the algorithm, the two available matches for forward rules are applied (in random order): *Leaf-FWD-Rule* translates the newly added *Class* `c2` and *Root-FWD-Rule* translates the *Package* `nRootP`; this results in the triple graph (c). Note that the last rule application creates a match for the repair rule *Connect-Root-Repair-Rule*. This is the reason why we start our synchronization process with applications of forward rules.

The incremental pattern matcher notifies about two broken consistency matches, which are dealt with in random order. `rootP` is no longer a root package (which is detected by a violation of the according filter NAC in the consistency pat-

tern) and `subP` is now a root package (which is detected by the missing incoming edge). Both violations are captured by repair rules, namely *Connect-Root-Repair-Rule* and *Make-Root-Repair-Rule*, whose applications lead to (d) and (e). The algorithm terminates with a triple graph that belongs to the TGG.

### 6.5 Prospect: support of further kinds of editing and advanced TGG features

We shortly describe the support of further kinds of editing and more advanced features of TGGs by our approach to synchronization, namely attributed TGGs, rules with NACs, and support for additional attribute constraints.

*Further kinds of editing* In our implementation (see Sect. 7), we do not only support the addition of elements and propagation of edits that correspond to source rules of derived edit rules. Actually, we do not make any assumptions about the kind of editing. This is achieved by incorporating the application of repair rules into the algorithm suggested by Leblebici et al. [41,42], which has also been proved to be correct and to terminate. The implemented algorithm first tries to apply a forward or repair rule. If there is none available with a valid match, the algorithm falls back to revoking of an invalid rule application. This means that all elements that have been created by this rule application are deleted (and adjacent edges of deleted nodes are implicitly deleted as well). In line with that revoking of invalid rule applica-

tions, it also allows for implicit deletion of adjacent edges in the application of repair rules. In that way, the application of a repair rule might trigger new appearances of broken consistency matches. We are convinced that *correctness* is not affected by that more general approach: Inspecting the proofs of Corollary 1 and Theorem 5, the key to correctness is that the sequences of currently valid consistency matches remain consistently marking. That is achieved via the conditions on matches for repair rules to be *reversing*, *context-preserving*, and *creation-preserving*. Dropping the condition to be *non-disabling* (by implicitly deleting adjacent edges) does not effect correctness, therefore. However, proving termination in that more general context is future work.

*Advanced features* The *attribution of graphs* can be formalized by representing data values as special nodes and the attribution of nodes and edges as special edges connecting graph elements with these data nodes [17]. As the rules of a TGG are monotonic, they only set attribute values but never delete or change them. (The deletion or change of an attribute value would include the deletion of the attribution edge pointing to it.) The formal construction of short-cut rules is based purely on category-theoretic concepts, which can be directly applied to rules on attributed triple graphs as well. The properties proven for short-cut rules in [22] are valid also in that case.<sup>8</sup> Hence, we can freely apply the construction of short-cut rules and derivation of repair rules to attributed TGGs. In fact, our implementation already supports attribution. For the propagation of attribute changes (made by a user), however, we rely on the inherent support eMoflon offers, which is discussed in Sect. 7. Deriving repair rules to propagate such changes is possible in principle but remains future work.

In practical applications, TGGs are often not only attributed but also equipped with *attribute constraints*. These enable the user to, for example, link the values of attributes of correlated nodes. eMoflon comes with facilities to detect violations of such constraints and offers support to repair such violations. In our implementation, we rely on these features of eMoflon to support attribute constraints but do not contribute additional support in our newly proposed synchronization algorithm.

To summarize, while fully formalized for the case of plain TGG rules without attribution, our implementation already supports the synchronization of attributed TGGs with additional attribute constraints. As these additional features do not affect our construction of short-cut and repair rules, we do not consider them (yet) to improve the propagation of attribute changes (that may lead to violations of attribute

constraints). Instead, we rely on the existing theory and facilities of eMoflon as introduced by Anjorin et al. [7]. In contrast, while computing short-cut and repair rules of rules with NACs is straightforward, adapting our synchronization algorithm to that case is future work and no tool support is available yet.

## 7 Implementation

Our implementation<sup>9</sup> of a model synchronizer using (short-cut) repair rules is built on top of the existing EMF-based, general-purpose graph and model transformation tool eMoflon [43,57,58]. eMoflon offers support for rule-based unidirectional and bidirectional graph transformations where the latter one uses TGGs. The model synchronizer implemented in eMoflon extends Algorithm 1 slightly. It allows any kind of user edit on the source part of a triple graph. If there are no forward or repair rules to fix a broken match, broken rule applications can be revoked. Revoking of rule applications has been the standard way of fixing broken matches. Hence, the implemented model synchronizer is a true extension of the previous synchronizer in eMoflon supporting the repair of broken applications.

In the following, we present the architecture behind our optimized model synchronizer first. Thereafter, we describe how the automatic calculation of short-cut and repair rules is implemented.

### 7.1 Tool architecture

Figure 17 depicts a UML component diagram to show the main components of eMoflon's bidirectional transformation engine. The architecture has two main components: *TGG Core* contains the core components of eMoflon and *Repair Framework* adds (short-cut) repair rules to eMoflon's functionality. The *TGG engine* manages the synchronization process and alters source, target, and correspondence model in order to restore consistency. For this purpose, it applies forward/ backward *operationalized TGG rules* to translate elements or revokes broken rule applications.

Finding matches in an incremental way is an important requirement for efficient model synchronization since minor model changes should be detectable without re-evaluating the whole model. For this reason, eMoflon relies on *incremental pattern matching* to detect the appearance of new matches as well as the disappearance of formerly detected ones. It uses different incremental pattern matchers such as Democles [55] and HiPE [1] and allows switching freely between them for optimizing the performance for each transformation

<sup>8</sup> To be precise, in [22], all proofs are elaborated for the case of monotonic rules in an *adhesive* category. Attributed triple graphs are *adhesive HLR* which is a weaker notion. However, inspecting the proofs, this does not make any difference as long as the category has so-called *effective pushouts*. This is known to be the case for attributed (triple) graphs; compare, e.g., [18, Remark 5.57].

<sup>9</sup> Both, the implementation and the evaluation, can be accessed via <https://github.com/Echtzeitsysteme/STTT-SC-Eval>.

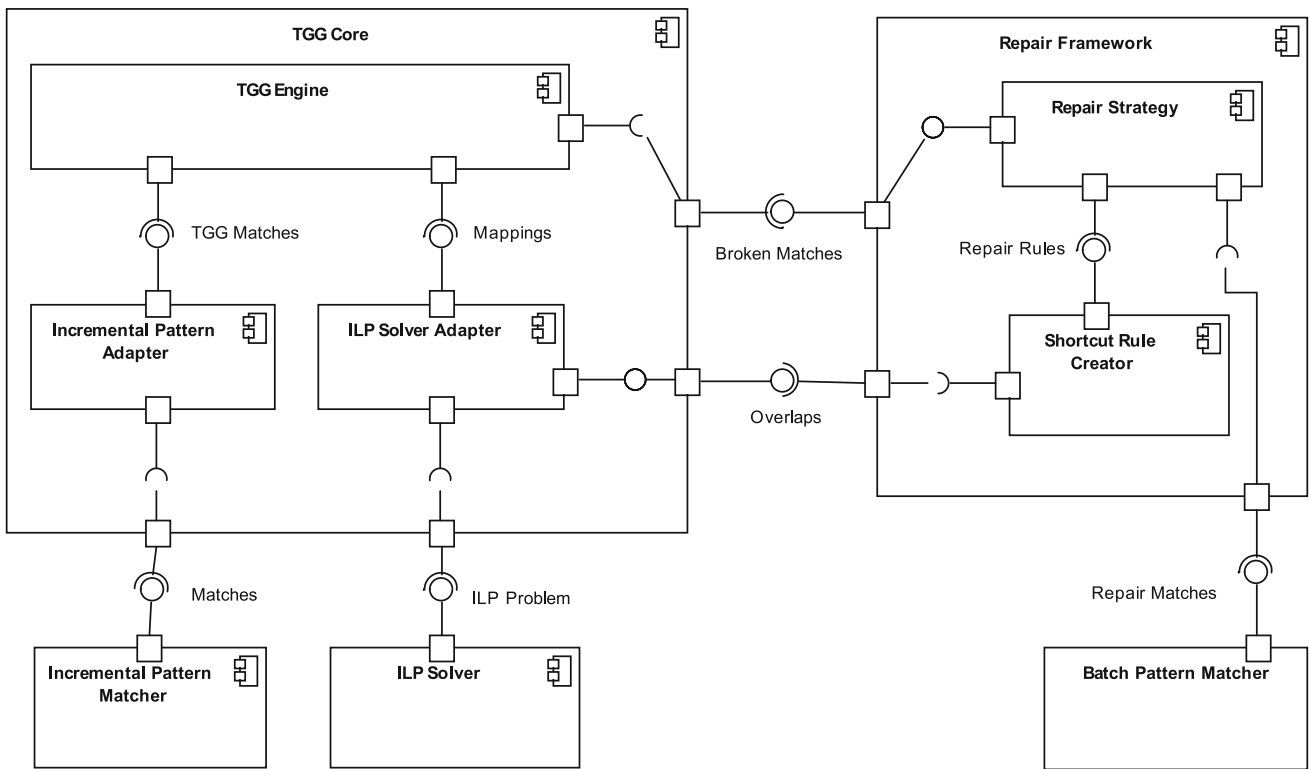


Fig. 17 eMoflon—Architecture of the bidirectional transformation engine

scenario. Furthermore, eMoflon employs the use of various *integer linear programming* (ILP) solvers such as Gurobi [28] and CPLEX [34], e.g., in order to find correspondence links (mappings) between source and target models, which is referred to as consistency check [46].

We have extended this basic setup by introducing the *Repair Framework*, which consists of the *Repair Strategy* and the *Shortcut Rule Creator*. The *Repair Strategy* is attached to the *TGG Engine* from which it is called with a set of broken rule matches. It attempts to repair the corresponding rule applications by using repair rules created by the *Shortcut Rule Creator*, which uses the ILP interface provided by the *TGG Core* in order to find overlaps between TGG rules and finally, to create short-cut repair rules. For invoking the repair rules, however, we have to find matches of repair rules. This is done by a *Batch (local-search) Pattern Matcher* which, in contrast to the incremental pattern matcher, does not perform any book-keeping. As a repair of a rule application is always done locally, the checking of matches throughout the whole model is considered to be too expensive and thus, a *Batch Pattern Matcher* can perform this task more efficiently.

### 7.2 ILP-based short-cut rule creation

In order to create an overlap between two rules, a morphism between the graphs of both rules has to be found: Each ele-

ment may only be mapped once; a context element may only be mapped to another context element. Created elements are mapped to each other, respectively. Furthermore, a node can only be mapped to a node of the same type as we do not incorporate inheritance between types yet. Edges are allowed to be mapped to each other only if their corresponding source and target elements are also mapped to each other, respectively.

We use integer linear programming (ILP) to encode the search space of all possible mappings and search for a maximal mapping. Each possible mapping  $m$  is considered to be a variable of our ILP problem such that calculating

$$\max(\sum_{m \in M} m)$$

yields the maximal overlap, with  $M$  being the set of all mappings and  $m \in \{0, 1\}$ . To ensure that each element  $e$  is mapped only once, we define a constraint to exclude non-used mappings:  $(\sum_{m \in A_e} m) \leq 1$  with  $A_e$  being the set of all alternative mappings for element  $e$ . To ensure that edges are mapped only if their adjacent nodes are mapped as well, we define the following constraint:  $m_e \implies m_v$  which translates to  $m_e \leq m_v$  with  $m_e$  being the edge mapping and  $m_v$  being one of the mappings of node  $src(e)$  or  $trg(e)$ . Maximizing the number of activated variables yields the common kernel of both input rules, i.e., a maximal overlap between them. If the overlap between the created elements of both

rules is empty, we drop this overlap as the resulting short-cut rule would not preserve any elements. Given a common kernel of two rules, we glue them along this kernel and yield a short-cut rule. For all elements of the resulting short-cut rule, which are not in the common kernel, we do the following: (1) Preserved elements remain preserved in the short-cut rule. (2) Created elements of the first rule become deleted ones as the first rule is inverted. (3) Created elements of the second rule remain created ones.

We calculate two kinds of overlap for each pair of rules and hence, two short-cut rules: a maximal and a minimal overlap. The maximal overlap is calculated by allowing mappings between all created and context elements, respectively. On the other hand, the minimal overlap is created by allowing mappings between created elements only. Considering the corresponding ILP problem, this means that all other mapping candidates are dropped.

Finally, the derived short-cut rules are operationalized to obtain the repair rules employed in our synchronization algorithm.

### 7.3 Attribute constraints

Although attribute constraints have not been incorporated formally in our approach, eMoflon is able to define and solve those within the former legacy translation and synchronization process. As can be seen in Fig. 19, many rules have an equality constraint defined between the name attributes of created elements on both, source and target parts. For TGG rules, this means that the attribute values may be chosen arbitrarily since both nodes would be created from scratch. In forward rules, source elements are already present which means that an attribute constraint can be interpreted as to propagate or copy the already present value to a newly created element. We reuse this functionality for our new synchronization process in the following way: After applying a repair rule, we ensure that the constraints of the replacing rule are fulfilled. The definition of attribute constraints and their treatment is due to Anjorin et al. [7].<sup>10</sup>

## 8 Evaluation

We evaluate our approach with respect to two aspects using the running example in an extended form. First, we investigate the performance of our approach w.r.t. information loss and execution time. A set of real and synthesized models is given which we use to apply four different kinds of model changes. Secondly, we evaluate the quality of our short-cut

rule generation strategy by comparing generated short-cut rules with well-known code refactorings.

Our experimental setup consists of 24 TGG rules (shown in “Appendix”) that specify consistency between Java AST and custom documentation models. In addition, there are 38 short-cut rules being derived from the set of TGG rules. A small modified excerpt of this rule set was given in Sect. 2. For this evaluation, however, we define consistency not only between *Package* and *Folder* hierarchies but also between type definitions, e.g., *Classes* and *Interfaces*, and *Fields* and *Methods* with their corresponding documentation entries.

### 8.1 Performance evaluation

To get realistic models, we extracted five models from Java projects hosted on Github using the reverse engineering tool MoDisco [12] and translated them into our own documentation structure. In addition, we generated five synthetic models consisting of  $n$ -level *Package* hierarchies with each non-leaf *Package* containing five sub-*Packages* and each leaf *Package* containing five *Classes*. While the realistic models shall show that our approach scales to real world cases, the synthetic models are chosen to show scalability in a more controlled way by increasing hierarchies gradually.

To evaluate our synchronization process, we performed several model changes. We refactored each of the models in four different scenarios; two example refactorings are the moving of a *Class* from one *Package* to another or the complete relocation of a *Package*. Then we used eMoflon to synchronize these changes in order to restore consistency to the documentation model using two synchronization processes, namely with and without *repair rules*. The legacy synchronization process of eMoflon is presented in [41,42]; the new synchronization process applying additional repair rules takes place according to the algorithm presented in Sect. 6 with the extensions mentioned in Sect. 6.5.

These synchronization steps are subject to our evaluation and we pose the following research questions: **(RQ1)** *For different kinds of model changes, how many elements can be preserved that would be deleted and recreated otherwise?* **(RQ2)** *How does our new synchronization process affect the runtime performance?* **(RQ3)** *Are there specific scenarios in which our new synchronization process performs especially good or bad?*

In the following, we evaluate our new *synchronization process by repair rules* against the *legacy synchronization process* in eMoflon. While the legacy one revokes forward rule applications and re-propagates the source model using forward rules, our new one prefers to apply short-cut repair rules as far as possible and falls back to revoking and re-propagation if there is no possible repair rule application. However, in our evaluation, it will not be necessary to apply revocation steps as we will see below.

<sup>10</sup> This approach allows one to specify constraints on attributes that involve also operations which are not only equality checks such as the concatenation of values of type *String*.



**Table 1** Legacy synchronizer—time in sec. and number of created elements

Models	Both		Legacy synchronization							
	Trans.		Scen. 1		Scen. 2		Scen. 3		Scen. 4	
	Sec.	Elts	Sec.	Elts	Sec.	Elts	Sec.	Elts	Sec.	Elts
tgg.core	6.4	1.6k	39	1.6k	3.8	99	0.64	17	0.2	3
modisco.java	9.9	3.2k	228	3.3k	18.6	192	3.6	33	0.4	4
eclipse.compare	10.74	3.8k	83	3.7k	3.1	76	2.36	47	0.1	1
eclipse.graphiti	20.7	6.5k	704	6.5k	63.9	490	5.65	25	0.9	3
synthetic $n = 1$	0.6	89	0.5	84	0.2	21	0.07	5	0.03	1
synthetic $n = 2$	1.4	345	1.7	340	0.2	21	0.11	5	0.04	1
synthetic $n = 3$	3.5	1369	13.2	1364	0.3	21	0.11	5	0.07	1
synthetic $n = 4$	14.5	5.5k	141.5	5.5k	1	21	0.32	5	0.09	1
synthetic $n = 5$	58.5	22k	2863	22k	10.7	21	1.07	5	0.23	1

To evaluate the performance of both legacy and the new model synchronization process, we consider four scenarios ranging from worst-case (Scenario 1) to best-case (Scenario 4) for the legacy implementation: Altering a root *Package* by creating a new *Package* as root would imply that all rule applications have to be reverted to synchronize the changes correctly with the legacy synchronization process (Scenario 1). In contrast, our new approach might perform poorly when a model change does not inflict a long cascade of invalid rule applications. Hence, we move *Classes* between *Packages* (Scenario 3) and *Methods* between *Classes* (Scenario 4) to measure if the effort of applying *repair rules* does infer a performance loss when both, the new and old algorithm, do not have to repair many broken rule applications. Note that Scenario 4 extends our evaluation presented in [23] as it provides a more fine-granular scenario. Finally, we simulate a scenario which is somewhat between the first three by relocating leaf *Packages* (Scenario 2) which, using the legacy model synchronization, would lead to a re-translation of all underlying elements.

Tables 1 and 2 depict the measured time in seconds (Sec.) and the number of re-/created elements (Elts) in each Scenario (1)–(4). The first table additionally shows measurements for the initial translation (Trans.) of the Java AST model into the documentation structure. For each scenario, Table 1 shows the numbers of synchronization steps using the legacy synchronizer without *repair rules* while Table 2 reflects the numbers of our new synchronizer with *repair rules*.

W.r.t. our research questions stated above, we interpret these tables as follows: The Elts columns of Table 2 show clearly that using repair rules preserves all those elements in our scenarios that are deleted and recreated by the legacy algorithm otherwise as shown in Table 1 (RQ1). The runtime shows a significant performance gain for Scenario 1 including a worst-case model change in which the legacy algorithm has to re-translate all elements (RQ2).

*Repair rules* do not introduce an overhead compared to the legacy algorithm as can be seen for the synthetic time measurements in Scenario 4 where only one rule application has to be repaired or reapplied (RQ2). Our new approach excels when the cascade of invalidated rule applications is long. Even if this is not the case, it does not introduce any measurable overhead compared to the legacy algorithm as shown in Scenarios 2, 3, and 4 (RQ3). Furthermore, the synthetic examples also show that the new synchronizer needs nearly constant time for synchronizing a model change, independent of the size of a model.

*Threats to validity* Our evaluation is based on five real world and five synthetic models. Of course, there exists a wide range of Java projects that differ significantly from each other w.r.t. their size, purpose, and developer style. Thus, the results may not be transferable to other projects. Nonetheless, we argue that the four larger models extracted from Github projects are representative since they are deduced from established tools of the Eclipse ecosystem. The synthetic models are also representative as they show the scalability of our approach in a more controlled environment with an increasing scaling factor. Together, realistic and synthetic models show that our approach does not only increase the performance of eMoflons synchronization process but also reduce the amount of re-created elements. Since each re-created element may contain information that would be lost during the process, we preserve this information and increase the overall quality of eMoflons synchronization results. In this evaluation, we selected four edit operations that are representative w.r.t. their dependency on other edit operations. In particular, only edits are considered that (i) correspond to source rules of short-cut rules we derive; (ii) consist of a single step. They may not be representative w.r.t. other aspects such as size or kind of change. We consider those aspects to be of minor importance in this context as dependency is the cause for deleting and recreating elements in the legacy synchronization process. Nonetheless, our implementation is also not

**Table 2** New synchronizer—time in sec. and number of created elements

Models	Synchronization by repair rules							
	Scen. 1		Scen. 2		Scen. 3		Scen. 4	
	Sec.	Elts	Sec.	Elts	Sec.	Elts	Sec.	Elts
lang.List	0.2	0	–	–	0.03	0	0.02	0
tgg.core	0.8	0	0.11	0	0.05	0	0.04	0
modisco.java	2.5	0	0.2	0	0.09	0	0.1	0
eclipse.compare	0.7	0	0.08	0	0.04	0	0.03	0
eclipse.graphiti	6.1	0	0.21	0	0.09	0	0.1	0
synthetic $n = 1$	0.1	0	0.05	0	0.03	0	0.05	0
synthetic $n = 2$	0.1	0	0.05	0	0.02	0	0.04	0
synthetic $n = 3$	0.1	0	0.07	0	0.02	0	0.04	0
synthetic $n = 4$	0.4	0	0.14	0	0.04	0	0.04	0
synthetic $n = 5$	1.5	0	0.37	0	0.09	0	0.06	0

**Table 3** An overview of short-cut based propagation of refactorings

Refactoring	Function of TGG rules/eMoflon	Function of short-cut rules
Extract Class	Creating Class & Reference as Field	Relocating Fields
Extract Superclass	Creating Superclass & deleting superfluous Fields & Methods	Relocating Methods & Fields
Inline Class	Deleting Class	Relocating Fields
Replace Superclass with Delegate	Creating Fields for Delegation	Revoke inheritance
Collapse Hierarchy	Deleting Sub/Super-Class & superfluous Methods/Fields	Relocating Methods & Fields
Introduce Parameter Object	Creating Class & deleting Parameters	
Change Function Declaration	Renaming Methods & Creating/Deleting/Renaming Parameters	
Move Method		Relocating Method
Combine Functions into Class	Deleting superfluous Parameter	Relocating Function/Method
Push-Up Method	Deleting superfluous Fields in Subclasses	Relocating Methods
Push-Down Method	Creating Fields in Subclasses	Relocating Methods
Remove Setting Method	Deleting Method	
Move Field		Relocating Fields
Push-Up Field	Deleting superfluous Fields in Subclasses	Relocating Fields
Push-Down Field	Creating Fields in Subclasses	Relocating Fields
Rename Field	Re-Evaluate Attribute Constraints	

able to derive short-cut rules that are able to handle multiple steps at once, due to our current heuristic of generating them. Hence, we plan to investigate this further in the near future. Finally, we limited our evaluation to one TGG rule set only but we experienced similar results for a broader range of TGGs from the eMoflon test zoo,<sup>11</sup> which also included more asymmetric TGGs.<sup>12</sup>

<sup>11</sup> Accessible via <https://github.com/eMoflon/emoflon-ibex-tests>.

<sup>12</sup> Such as TerraceHouses2BlockSet, which is also accessible via the eMoflon test zoo.

## 8.2 Refactorings

As explained in Sect. 7, we currently employ two different strategies to overlap two rules and to create a short-cut rule. We pose the following research question: **(RQ4)** *Are the generated short-cut rules applicable to realistic scenarios? Are further short-cut rules necessary?* Since our example addresses code changes that are incorporated by the Java AST model primarily, we relate our approach to available code refactorings. In the following, we refer to the book on code refactorings written by Martin Fowler [21] which presents 66 refactorings.

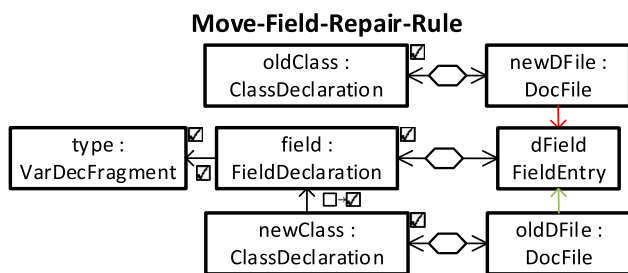


Fig. 18 Move-Field-Repair-Rule

Our example TGG, depicted in Fig. 19, defines consistency on a structural level solely, without incorporating behavior, i.e., the bodies of methods and constructors. Hence, we selected those refactorings that describe changes on *Packages*, *Classes* and *Interfaces*, *MethodDeclarations* and *Parameters*, and *Fields*. The result is a set of 16 refactorings for which we evaluated if short-cut rules help to directly propagate the corresponding change of the AST model or deletion and recreation has to take place.

Table 3 lists these refactorings together with information on the TGG rules and/or short-cut rules that are applicable in these scenarios. For some of the refactorings as e.g., *Extract Class* and *Push-Down Field*, we identified situations where not only short-cut rules are necessary to propagate the changes. In these cases, new elements may be created which can be propagated using operationalized TGG rules. The deletion of elements can be propagated by revoking the corresponding prior propagation step. For the reevaluation of attributes (e.g., for the refactoring *Rename Field*), we rely on the facilities of eMoflon. However, many refactorings benefit from using short-cut rules, for example, those that move methods and fields. If recreation of documentation on the target part is necessary, it can lead to information loss as there may not be all the necessary information in the Java AST model.

**Example 10** *Push-Up Field* moves and merges a similar field from various subclasses into a common superclass. If one of the subclass fields is moved to the superclass, we can propagate this change using *Move-Field-Repair-Rule*, which is depicted in Fig. 18.

In summary, we are able to solve all 16 refactorings using a combination of (inverse) TGG rules and our generated short-cut rules (RQ4).

*Threats to validity* Note that short-cut rules are especially useful when elements are moved instead of deleting and recreating them in some other location. Those changes are hard to detect and are not covered here. Refactorings such as *Push-Up Method*, which moves a method that occurs in several subclasses to their common superclass, can be done in two different ways. First, one of the methods is moved to the superclass while the methods in the other subclasses are

deleted. This employs the use of short-cut rules for the moved method followed by revocation steps for the deleted methods to delete the corresponding documentation elements. Second, all methods may be deleted and a new similar method is created in the superclass. In that case, there is no short-cut rule that helps to preserve information and all propagated documentation elements for the method will be blank. Hence, our approach depends on the kind of change. In particular, it helps when user edits also try to preserve information instead of recreating them.

In addition, we have not incorporated behavior in our example; such an extension of our TGG may be considered in future work. However, we can argue that most of those refactorings can be reduced to the movement of elements, the deletion of superfluous elements and the creation of new elements. These changes are manageable in general using a sequence of short-cut rule and (inverse) operationalized TGG rule applications.

Finally, we evaluated these cases by hand based on the generated short-cut rules from our implementation. Nonetheless, test cases implementing the identified refactorings and combinations of them will be made accessible via eMoflons test zoo. We hope that these tests will suite as a base to compare different sequential synchronizers in the future.

## 9 Related work

In this section, we relate our new model synchronization approach to already existing incremental model synchronization approaches. First, we discuss other TGG-based approaches in detail before relating to other bidirectional transformation (bx) approaches; these are considered more roughly. Finally, we mention some unidirectional approaches that are closely related to incremental model transformation and model repair. Work that is related to our use of partial triple graphs but not to model synchronization is considered in [37].

*TGG-based approaches to incremental model synchronization* Synchronization approaches are supposed to comply with the least-change property, which means that no unnecessary deletions and thus information loss should take place while restoring consistency. An overview of TGG-based least-change synchronization has been given by Stojkovic et al. [52]. The first part of our related work is based on that presentation.

Several approaches to model synchronization based on TGGs suffer from the fact that the revocation of a rule application may trigger the revocation of all dependent rule applications as well [26,40–42]. Such cascades of deletions shall be avoided to decrease runtime and unnecessary information loss.

Leveraging an incremental pattern matcher for TGG-based model synchronization was first suggested in [41,42]. Proofs of termination, correctness, and completeness are given. Moreover, the approach is implemented. In fact, this is the legacy synchronization we evaluated against in Sect. 8. As already mentioned, that approach revokes invalid consistency matches as long as there are any and subsequently, applies forward rules to translate yet untranslated elements. So, that approach is a typical example where a lot of unnecessary deletions may take place.

Hermann et al. [30] proposed a synchronization algorithm where, after an edit on the source part, first those correspondence elements are deleted that do not refer to an element in the source graph any longer. Thereafter, they parse the remaining triple graph to find the maximal, still valid sub-model. This model is used as a starting point to propagate the remaining changes from source to correspondence and target graphs using forward rules. The approach is completely formalized and proven to be correct, also for attributed TGGs; it can be applied to TGGs with deterministic<sup>13</sup> sets of operationalized rules. That approach avoids some unnecessary deletions but there are some that still can occur. In fact, the amount of unnecessary deletion taking place in that approach is dependent on the given TGG rules; a concrete example for that is given in [52]. While that approach is definitely a valuable contribution towards least-change synchronization, repeated parsing for maximally consistent sub-models is highly inefficient and might not scale to large models. At least part of that approach is implemented as HENSINTGG [20] using AGG [53] to perform necessary dependency checks on derived rules. As that approach focusses on correctness, completeness, and invertibility, the amount of achieved incrementality as well as principles of least change are not discussed in [30].

In [24], Giese and Hildebrandt propose rules that save nodes instead of deleting and re-creating them. In particular, they present a rule that directly propagates the movement of elements, i.e., the redirection of edges between existing elements. Moreover, they suggest to try a re-use of elements before deleting them. But they neither present a general construction for their rules nor formalize the re-use that takes place. Consequently, no proof of correctness is given. Instead, it is left as future work in [25]. The additional propagation rules that are given exemplary in [24] can be automatically derived as *repair rules* using our approach. In [10], Blouin et al. also add specifically designed repair rules to the rule set of their case study for avoiding information loss. Those example rules can be realized as *repair rule* in our approach as well.

In a similar vein, Greenyer et al. [27] propose to delete elements not directly but to mark them for deletion and to allow for their re-use in rule applications during synchronization. Only elements that cannot be re-used are deleted at the very end of synchronization. But that approach comes without any formalization and proof of correctness as well.

In contrast, the idea of re-using elements in model synchronizations has been rigorously formalized by Orejas and Pino [50]. They introduced *forward translation rules with reuse* and proposed a synchronization algorithm based on those rules. That algorithm is actually proven to be correct; moreover, it is incremental (in a technical sense). The practical effects of applying a *repair rule* in our approach and in their approach are very similar. While our *repair rules* allow for reuse and perform necessary deletions on the correspondence and target parts directly, their forward translation rules allow for a reuse where necessary deletions are performed at the end of a synchronization in a separate step. They need some additional technical infrastructure to determine the exact amount of necessary deletion. To the best of our knowledge, their approach has not been implemented yet.

In a guideline on how to develop a TGG, Anjorin et al. [5] explain how certain kinds of rules in a TGG avoid the loss of information better than others. There is empirical evidence that, following these guidelines, synchronization can be considerably accelerated compared to a batch mode as long as there is no need for additional offline recognition of model differences [45]. Transforming a given TGG into that form, however, may change the defined language and thus, is not always applicable. For example, the grammar of our running example allows generating hierarchies of *Packages* that constitute a set of disconnected trees. For meeting the suggestions in [5], a naive change of this grammar may change the language such that arbitrary graphs can be generated. That effect can be avoided by, e.g., designing suitable NACs for the rules and proving the equality of the generated model languages. That effort is not needed when following our approach.

In summary, it is well-known in the literature that there are a lot of situations where the derived forward rules of a TGG (and the revocation of their applications) are not suitable to efficiently propagate changes from source to target models. Several formal and informal approaches have been suggested to avoid this problem, at least partly. Table 4 provides an overview of all the approaches described above. It indicates the degree of information loss and presents whether the approach is automated, whether correctness of the proposed synchronization algorithm is proven, whether it has been (prototypically) implemented, and whether any performance gain could be shown for it. Our approach is based on the automated derivation of *repair rules*; it is able to comply with all the above categories. The correctness has been shown for model synchronization with repair rules. As our

<sup>13</sup> Deterministic in the sense that there are no competing rules for any translated element.



**Table 4** An overview of TGG-based synchronization approaches

	Degree of information loss	Automated	Correctness proven	Implemented	Evaluated performance gain
[41,42]	High	Yes	Yes	Yes	Yes
[26]	High	Yes	Only partially in [25]	Yes	Yes
[40]	High	Yes	Yes	Yes	Yes
[30]	To some extent	Yes	Yes	At least partially	No
[24]	Low	Yes	No	Yes	Yes
[27]	Low	Yes	No	Yes	No
[50]	Low	Yes	Yes	No	No
[5,45]	Low	Not needed	Not needed	Yes	Yes
[48,54]	Low	No	Yes	Yes	No
Ours	Low	Yes	Yes	Yes	Yes

implemented synchronization process can also revoke forward rules, the correctness proof has to be slightly extended to cover also that case which seems to be straight forward (see discussion in Sect. 6.5). Furthermore, support for some additional features of TGGs like NACs and attribution is future work (NACs) or not rigorously formalized (attribution).

*Comparison to other bx approaches* Anjorin et al. [4] compared three state-of-the-art bx tools, namely eMoflon [43] (rule-based), mediniQVT [2] (constraint-based), and BiGUL [36] (bx programming language) w.r.t. model synchronization. They point out that synchronization with eMoflon is faster than with both other tools as the runtimes of those tools all correlate with the overall model size while the runtime of eMoflon correlates with the size of the changes done by edit operations. Furthermore, eMoflon is the only tool that was able to solve all but one synchronization scenario while mediniQVT failed in four and BiGUL in two scenarios. One scenario was not solved because the solution with eMoflon deletes more model elements than absolutely necessary in that case. Using short-cut repair rules, we can solve the remaining scenario and moreover, can further increase the performance of eMoflon when solving model synchronization tasks. Macedo and Cunha present bidirectional model transformations based on ATL in [47]. By using the SAT solver Alloy, they are able to guarantee least-change model synchronization where two metrics are supported measuring change: the *graph edit distance* and the *operation-based distance*. While the synchronization results may be very good, this solver-based approach does not scale for large models. All this suggests that our tool is highly competitive, not only among TGG-based tools but also in comparison to other bx tools.

With regard to theoretical considerations, least change and incremental synchronization have also been actively investigated in other approaches, in particular when using lenses, e.g., [15,31–33,56]. The approach by Wang et al. [56] seems to be the most similar one to ours. That approach derives

functions to directly propagate changes from a source to a view and is applicable to tree-shaped data structures. As those approaches are less close to our work, detailed formal comparisons are left to future work.

*Further related works* *Change-preserving model repair* as presented in [48,54] is closely related to our approach. Assuming a set of consistency-preserving rules and a set of edit rules to be given, each edit rule is accompanied by one or more repair rules completing the edit step if possible. Such a complement rule is considered as repair rule of an edit rule w.r.t. an overarching consistency-preserving rule. Operationalized TGG rules fit into that approach but provide more structure: As graphs and rules are structured in triples, a source rule is also an edit rule being complemented by a forward rule. In contrast to that approach, source and forward rules can be automatically deduced from a given TGG rule. By our use of short-cut rules, we introduce a pre-processing step to first enlarge the sets of consistency-preserving rules and edit rules. Furthermore, the repair process presented in that paper has more restrictive presumptions than our synchronization process using repair rules w.r.t. independence of rule applications.

Boronat [11] presents an incremental uni-directional transformation approach. When retranslating a model after a change, affected elements of the old model are marked first and then, if possible, re-used instead of deleted and re-created (similar to the approaches suggested in [27,50] for TGGs). Again, the same effects can be obtained by constructing and applying short-cut rules but there, for plain graph transformation. A correctness proof for that approach is still missing.

## 10 Conclusion

Model synchronization, i.e., the task of restoring the consistency between two models after model changes, poses challenges to modern bidirectional model transformation

approaches and tools: We expect them to synchronize changes without unnecessary loss of information and to show a reasonable performance. Here, we restrict ourselves to model synchronizations where only one model is changed at a time.

While Triple Graph Grammars (TGGs) provide the means to perform model synchronization tasks in general, efficient model synchronization without unnecessary information loss may not always be fulfilled since basic TGG rules are not designed to support intermediate model editing and repair. Therefore, we propose to add short-cut rules, a special form of generalized TGG rules that allow taking back one edit action and to perform an alternative one. In our evaluation, we show that repair rules derived from short-cut rules allow for a kind of incremental model synchronization with considerably decreased information loss and improved runtime compared to synchronization without these rules.

In this paper, we show the correctness of our synchronization approach, present the implementation design, and evaluate the corresponding tool support w.r.t. performance and unnecessary information loss. While the tool support already covers attributes of model elements, the correctness proof of our synchronization approach w.r.t. to these extensions is prepared but still up to future work.

While model synchronization means the propagation of model changes from one view to another, model changes may also occur concurrently on both views of a model. Hence, model synchronization approaches have to cover those scenarios as well. Short-cut rules may also be promising to avoid information loss in that more general setting; they have not been considered in the context of other approaches to concurrent model synchronization in the literature [49,60]. As changes of both model views may be in conflict with each other, the development of an efficient concurrent model synchronization process which avoids unnecessary information loss poses a challenge for future work.

**Acknowledgements** This work was partially funded by the German Research Foundation (DFG) project “Triple Graph Grammars (TGG) 2.0”. We would also like to thank the anonymous reviewers for their thoughtful comments and efforts.

**Funding** Open Access funding provided by Projekt DEAL.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## Appendix: A evaluation ruleset

In this section, we present additional information related to our evaluation from Sect. 8.

Figure 19 depicts the full TGG rule set used of our evaluation. The first rule *JavaModel-2-DocModel-Rule* defines consistency between a MoDisco *Model* and a *DocModel* that contains three sub *DocModels* and another *Folder* linked to the common *DocModel*. These different containers are used to separate Java entities on the documentation site to split them up into common Java data types, external Java references and source references. *JavaModel-2-DocModel-Rule* then defines consistency between *Packages* and *Folder* given that their parent are a MoDisco *Model* and a *DocModel*, respectively. Using *JavaPackage-2-DocFolder-Rule*, we can now create *Package* and *Folder* hierarchies recursively. Furthermore, there are four rules that define consistency for *ClassesDeclarations*, *InterfacesDeclaration*, *EnumDeclaration* and inner *ClassesDeclarations* each with a *Doc-File*. Also, for the nine primitive types, e.g., boolean, byte and short, consistency is defined between each of them and a *Doc-File*. Given a *ClassDeclaration* or an *InterfaceDeclaration* with its corresponding *Doc-File*, we also define consistency between *MethodDeclarations* on one and *MethodEntries* on the other side. Using the consistency between methods on both sides, we are able to define consistency between *TypeAccesses* and *Parameters*, once for method signatures and once for the return statement. Finally, we define consistency between generalization and realization relationships using three rules. First, a rule for *ClassesDeclarations* that extend another *ClassDeclaration*, second a rule for *InterfacesDeclaration* extending another *InterfaceDeclaration* and last for *ClassesDeclarations* implementing an *InterfaceDeclaration*.

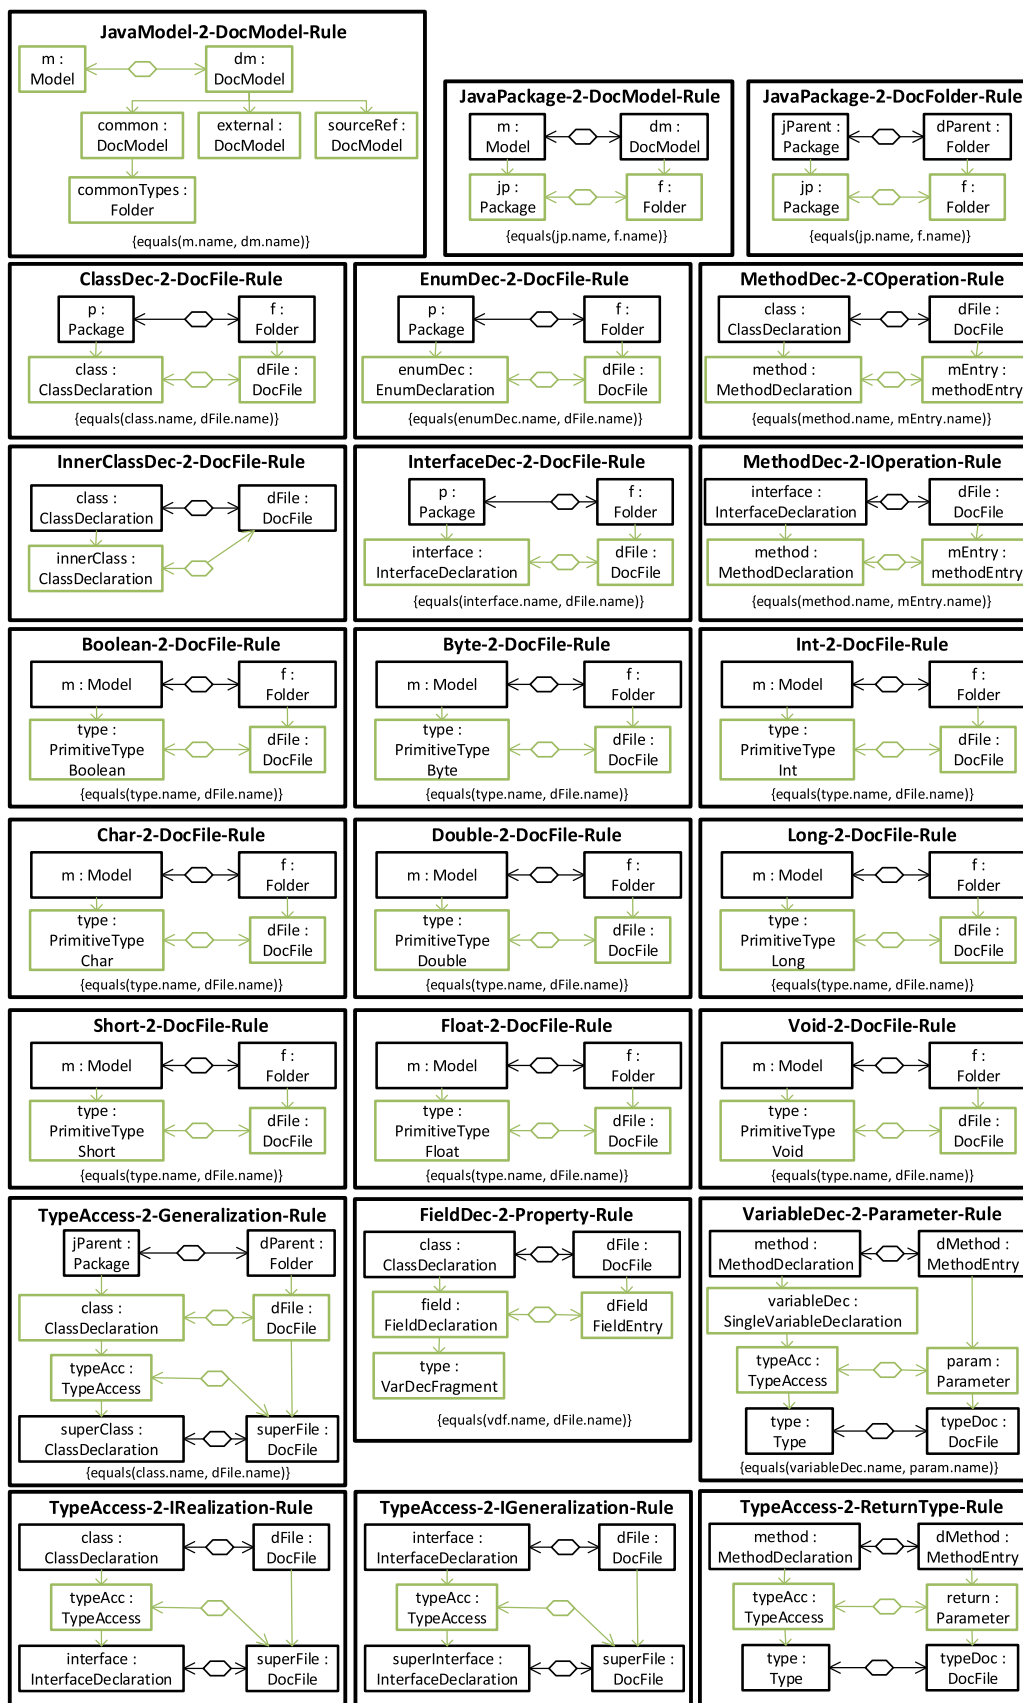


Fig. 19 Evaluation—TGG rule set

## References

1. HiPE: Highly (Scalable) Incremental Pattern matching Engine. <https://github.com/Arikae/HiPE-Updatesite>. Accessed 27 Aug 2020
2. Ikv++: Medini QVT. <http://projects.ikv.de/qvt>. Accessed 27 Aug 2020
3. Abou-Saleh, F., Cheney, J., Gibbons, J., McKinna, J., Stevens, P.: Introduction to bidirectional transformations. In: Gibbons, J., Stevens, P. (eds.) *Bidirectional Transformations—International Summer School, Lecture Notes in Computer Science*, vol. 9715, pp. 1–28. Springer, Berlin (2018)
4. Anjorin, A., Diskin, Z., Jouault, F., Ko, H., Leblebici, E., Westfechtel, B.: Benchmark reloaded: a practical benchmark framework for bidirectional transformations. In: *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.*, pp. 15–30 (2017). <http://ceur-ws.org/Vol-1827/paper6.pdf>
5. Anjorin, A., Leblebici, E., Kluge, R., Schürr, A., Stevens, P.: A systematic approach and guidelines to developing a triple graph grammar. In: Cunha, A., Kindler, E. (eds.) *Proceedings of the 4th International Workshop on Bidirectional Transformations co-located with Software Technologies: Applications and Foundations (STAF 2015), CEUR Workshop Proceedings*, vol. 1396, pp. 66–80. CEUR-WS.org (2015). <http://ceur-ws.org/Vol-1396>
6. Anjorin, A., Leblebici, E., Schürr, A., Taentzer, G.: A static analysis of non-confluent triple graph grammars for efficient model transformation. In: Giese, H., König, B. (eds.) *Graph Transformation—7th International Conference, ICGT 2014, Held as Part of STAF 2014, York, UK, July 22–24, 2014. Proceedings, Lecture Notes in Computer Science*, vol. 8571, pp. 130–145. Springer (2014). [https://doi.org/10.1007/978-3-319-09108-2\\_9](https://doi.org/10.1007/978-3-319-09108-2_9)
7. Anjorin, A., Varró, G., Schürr, A.: Complex attribute manipulation in tggs with constraint-based programming techniques. *Electron. Commun. EASST* (2012). <https://doi.org/10.14279/tuj.eceasst.49.707>
8. Baldan, P., Corradini, A., Heindel, T., König, B., Sobociński, P.: Processes and unfoldings: concurrent computations in adhesive categories. *Math. Struct. Comput. Sci.* (2014). <https://doi.org/10.1017/S096012951200031X>
9. Biermann, E., Ermel, C., Taentzer, G.: Formal foundation of consistent EMF model transformations by algebraic graph transformation. *Softw. Syst. Model.* **11**(2), 227–250 (2012)
10. Blouin, D., Plantec, A., Dissaux, P., Singhoff, F., Diguët, J.P.: Synchronization of models of rich languages with triple graph grammars: an experience report. In: Di Ruscio, D., Varró, D. (eds.) *Theory and Practice of Model Transformations*, pp. 106–121. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-08789-4\\_8](https://doi.org/10.1007/978-3-319-08789-4_8)
11. Boronat, A.: Offline delta-driven model transformation with dependency injection. In: Hähnle, R., van der Aalst, W. (eds.) *Fundamental Approaches to Software Engineering*, pp. 134–150. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-16722-6\\_8](https://doi.org/10.1007/978-3-030-16722-6_8)
12. Brunelière, H., Cabot, J., Dupé, G., Madiot, F.: Modisco: a model driven reverse engineering framework. *Inf. Softw. Technol.* **56**(8), 1012–1032 (2014). <https://doi.org/10.1016/j.infsof.2014.04.007>
13. Cheney, J., Gibbons, J., McKinna, J., Stevens, P.: On principles of least change and least surprise for bidirectional transformations. *J. Object Technol.* **16**(1), 3:1–3:31 (2017). <https://doi.org/10.5381/jot.2017.16.1.a3>
14. Czarnecki, K., Foster, J.N., Hu, Z., Lämmel, R., Schürr, A., Terwilliger, J.F.: Bidirectional transformations: a cross-discipline perspective. In: Paige, R.F. (ed.) *Theory and Practice of Model Transformations*, pp. 260–283. Springer, Berlin (2009)
15. Diskin, Z., Xiong, Y., Czarnecki, K.: From state- to delta-based bidirectional model transformations: the asymmetric case. *J. Object Technol.* **10**(6), 1–25 (2011). <https://doi.org/10.5381/jot.2011.10.1.a6>
16. Ehrig, H., Ehrig, K., Ermel, C., Hermann, F., Taentzer, G.: Information preserving bidirectional model transformations. In: Dwyer, M.B., Lopes, A. (eds.) *Fundamental Approaches to Software Engineering*, pp. 72–86. Springer, Berlin (2007). [https://doi.org/10.1007/978-3-540-71289-3\\_7](https://doi.org/10.1007/978-3-540-71289-3_7)
17. Ehrig, H., Ehrig, K., Prange, U., Taentzer, G.: *Fundamentals of Algebraic Graph Transformation. Monographs in Theoretical Computer Science*. Springer, Berlin (2006)
18. Ehrig, H., Ermel, C., Golas, U., Hermann, F.: *Graph and Model Transformation—General Framework and Applications. Monographs in Theoretical Computer Science. An EATCS Series*. Springer, Berlin (2015). <https://doi.org/10.1007/978-3-662-47980-3>
19. Ehrig, H., Golas, U., Habel, A., Lambers, L., Orejas, F.:  *$\mathcal{M}$ -adhesive transformation systems with nested application conditions. Part I: parallelism, concurrency and amalgamation. Math. Struct. Comput. Sci.* **24**(4) (2014)
20. Ermel, C., Hermann, F., Gall, J., Binanzer, D.: Visual modeling and analysis of EMF model transformations based on triple graph grammars. *ECEASST* (2012). <https://doi.org/10.14279/tuj.eceasst.54.771>
21. Fowler, M.: *Refactoring: Improving the Design of Existing Code*, 2nd edn. Addison-Wesley, Boston (2018)
22. Fritsche, L., Kosiol, J., Schürr, A., Taentzer, G.: Short-cut rules. Sequential composition of rules avoiding unnecessary deletions. In: Mazzara, M., Ober, I., Salaün, G. (eds.) *Software Technologies: Applications and Foundations*, pp. 415–430. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-04771-9\\_30](https://doi.org/10.1007/978-3-030-04771-9_30)
23. Fritsche, L., Kosiol, J., Schürr, A., Taentzer, G.: Efficient model synchronization by automatically constructed repair processes. In: Hähnle, R., van der Aalst, W. (eds.) *Fundamental Approaches to Software Engineering*, pp. 116–133. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-16722-6\\_7](https://doi.org/10.1007/978-3-030-16722-6_7)
24. Giese, H., Hildebrandt, S.: Efficient model synchronization of large-scale models. *Tech. Rep. 28*, Hasso-Plattner-Institut (2009)
25. Giese, H., Hildebrandt, S., Lambers, L.: Bridging the gap between formal semantics and implementation of triple graph grammars. *Softw. Syst. Model.* **13**(1), 273–299 (2014). <https://doi.org/10.1007/s10270-012-0247-y>
26. Giese, H., Wagner, R.: From model transformation to incremental bidirectional model synchronization. *Softw. Syst. Model.* **8**(1), 21–43 (2009). <https://doi.org/10.1007/s10270-008-0089-9>
27. Greenyer, J., Pook, S., Rieke, J.: Preventing information loss in incremental model synchronization by reusing elements. In: France, R.B., Kuester, J.M., Bordbar, B., Paige, R.F. (eds.) *Modelling Foundations and Applications. Proceedings of the 7th European Conference on Modelling Foundations and Applications*, pp. 144–159. Springer, Berlin (2011). [https://doi.org/10.1007/978-3-642-21470-7\\_11](https://doi.org/10.1007/978-3-642-21470-7_11)
28. Gurobi Optimization, I.: *Gurobi optimizer reference manual* (2017). [https://www.gurobi.com/wp-content/plugins/hd\\_documentations/documentation/7.0/refman.pdf](https://www.gurobi.com/wp-content/plugins/hd_documentations/documentation/7.0/refman.pdf). Accessed 27 Aug 2020
29. Hermann, F., Ehrig, H., Golas, U., Orejas, F.: Efficient analysis and execution of correct and complete model transformations based on triple graph grammars. In: *Proceedings of the First International Workshop on Model-Driven Interoperability, MDI '10*, pp. 22–31. ACM, New York, NY, USA (2010) <https://doi.org/10.1145/1866272.1866277>
30. Hermann, F., Ehrig, H., Orejas, F., Czarnecki, K., Diskin, Z., Xiong, Y., Gottmann, S., Engel, T.: Model synchronization based on triple graph grammars: correctness, completeness and invert-



- ibility. *Softw. Syst. Model.* **14**(1), 241–269 (2015). <https://doi.org/10.1007/s10270-012-0309-1>
31. Hinkel, G., Burger, E.: Change propagation and bidirectionality in internal transformation DSLs. *Softw. Syst. Model.* **18**(1), 249–278 (2019). <https://doi.org/10.1007/s10270-017-0617-6>
  32. Hofmann, M., Pierce, B., Wagner, D.: Edit Lenses. In: *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '12*, pp. 495–508. ACM, New York, NY, USA (2012). <https://doi.org/10.1145/2103656.2103715>
  33. Horn, R., Perera, R., Cheney, J.: Incremental relational lenses. *Proc. ACM Program. Lang.* **2**(ICFP), 1–30 (2018). <https://doi.org/10.1145/3236769>
  34. IBM ILOG CPLEX Optimization Studio CPLEX User's Manual (Version 12 Release 8). [https://www.ibm.com/support/knowledgecenter/SSSA5P\\_12.8.0/ilog.odms.studio.help/pdf/usrcplex.pdf](https://www.ibm.com/support/knowledgecenter/SSSA5P_12.8.0/ilog.odms.studio.help/pdf/usrcplex.pdf). Accessed 27 Aug 2020
  35. Klar, F., Lauder, M., Königs, A., Schürr, A.: Extended triple graph grammars with efficient and compatible graph translators. In: Engels, G., Lewerentz, C., Schäfer, W., Schürr, A., Westfechtel, B. (eds.) *Graph Transformations and Model-Driven Engineering: Essays Dedicated to Manfred Nagl on the Occasion of his 65th Birthday*, pp. 141–174. Springer, Berlin (2010). [https://doi.org/10.1007/978-3-642-17322-6\\_8](https://doi.org/10.1007/978-3-642-17322-6_8)
  36. Ko, H., Zan, T., Hu, Z.: Bigul: a formally verified core language for putback-based bidirectional programming. In: *Proceedings of the 2016 ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation, PEPM 2016*, St. Petersburg, FL, USA, January 20–22, 2016, pp. 61–72 (2016). <https://doi.org/10.1145/2847538.2847544>
  37. Kosiol, J., Fritsche, L., Schürr, A., Taentzer, G.: Adhesive subcategories of functor categories with instantiation to partial triple graphs. In: Guerra, E., Orejas, F. (eds.) *Graph Transformation*, pp. 38–54. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-23611-3\\_3](https://doi.org/10.1007/978-3-030-23611-3_3)
  38. Kreowski, H.: Is parallelism already concurrency? Part 1: Derivations in graph grammars. In: Ehrig, H., Nagl, M., Rozenberg, G., Rosenfeld, A. (eds.) *Graph-Grammars and Their Application to Computer Science, 3rd International Workshop, Warrenton, Virginia, USA, December 2–6, 1986, Lecture Notes in Computer Science*, vol. 291, pp. 343–360. Springer (1986). [https://doi.org/10.1007/3-540-18771-5\\_63](https://doi.org/10.1007/3-540-18771-5_63)
  39. Lack, S., Sobociński, P.: Adhesive and quasiadhesive categories. *Theor. Inf. Appl.* **39**(3), 511–545 (2005). <https://doi.org/10.1051/ita:2005028>
  40. Lauder, M., Anjorin, A., Varró, G., Schürr, A.: Efficient model synchronization with precedence triple graph grammars. In: Ehrig, H., Engels, G., Kreowski, H.J., Rozenberg, G. (eds.) *Graph Transformations*, pp. 401–415. Springer, Berlin (2012). [https://doi.org/10.1007/978-3-642-33654-6\\_27](https://doi.org/10.1007/978-3-642-33654-6_27)
  41. Leblebici, E.: Inter-model consistency checking and restoration with triple graph grammars. Ph.D. thesis, Darmstadt University of Technology, Germany. <http://tuprints.ulb.tu-darmstadt.de/7426/> (2018). Accessed 27 Aug 2020
  42. Leblebici, E., Anjorin, A., Fritsche, L., Varró, G., Schürr, A.: Leveraging incremental pattern matching techniques for model synchronisation. In: de Lara, J., Plump, D. (eds.) *Graph Transformation*, pp. 179–195. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-61470-0\\_11](https://doi.org/10.1007/978-3-319-61470-0_11)
  43. Leblebici, E., Anjorin, A., Schürr, A.: Developing eMoflon with eMoflon. In: Di Ruscio, D., Varró, D. (eds.) *Theory and Practice of Model Transformations*, pp. 138–145. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-08789-4\\_10](https://doi.org/10.1007/978-3-319-08789-4_10)
  44. Leblebici, E., Anjorin, A., Schürr, A., Taentzer, G.: Multi-amalgamated triple graph grammars: formal foundation and application to visual language translation. *J. Vis. Lang. Comput.* **42**, 99–121 (2017). <https://doi.org/10.1016/j.jvlc.2016.03.001>
  45. Leblebici, E., Anjorin, A., Schürr, A.: A catalogue of optimization techniques for triple graph grammars. In: Fill, H., Karagiannis, D., Reimer, U. (eds.) *Modellierung 2014*, 19.–21. März 2014, Wien, Österreich, LNI, vol. 225, pp. 225–240. GI. <https://dl.gi.de/handle/20.500.12116/20941> (2014). Accessed 27 Aug 2020
  46. Leblebici, E., Anjorin, A., Schürr, A.: Inter-model consistency checking using triple graph grammars and linear optimization techniques. In: *International Conference on Fundamental Approaches to Software Engineering*, pp. 191–207. Springer (2017)
  47. Macedo, N., Cunha, A.: Least-change bidirectional model transformation with QVT-R and ATL. *Softw. Syst. Model.* **15**(3), 783–810 (2016). <https://doi.org/10.1007/s10270-014-0437-x>
  48. Ohrndorf, M., Pietsch, C., Kelter, U., Kehrer, T.: Revision: a tool for history-based model repair recommendations. In: *Proceedings of the 40th International Conference on Software Engineering: Companion Proceedings, ICSE 2018, Gothenburg, Sweden, May 27–June 03, 2018*, pp. 105–108. ACM (2018). <https://doi.org/10.1145/3183440.3183498>
  49. Orejas, F., Boronat, A., Ehrig, H., Hermann, F., Schölzel, H.: On propagation-based concurrent model synchronization. *ECEASST* (2013). <https://doi.org/10.14279/tuj.eceasst.57.871>
  50. Orejas, F., Pino, E.: Correctness of incremental model synchronization with triple graph grammars. In: Di Ruscio, D., Varró, D. (eds.) *Theory and Practice of Model Transformations*, pp. 74–90. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-08789-4\\_6](https://doi.org/10.1007/978-3-319-08789-4_6)
  51. Schürr, A.: Specification of graph translators with triple graph grammars. In: Mayr, E.W., Schmidt, G., Tinhofer, G. (eds.) *Graph-Theoretic Concepts in Computer Science, Lecture Notes in Computer Science*, vol. 903, pp. 151–163. Springer, Berlin (1995). [https://doi.org/10.1007/3-540-59071-4\\_45](https://doi.org/10.1007/3-540-59071-4_45)
  52. Stojkovic, M., Laux, S., Anjorin, A.: Existing and new ideas on least change triple graph grammars. 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. 1–5. CEUR-WS.org (2017). <http://ceur-ws.org/Vol-1827/paper2.pdf>
  53. Taentzer, G.: AGG: A graph transformation environment for modeling and validation of software. In: Pfaltz, J.L., Nagl, M., Böhlen, B. (eds.) *Applications of Graph Transformations with Industrial Relevance, Second International Workshop, AGTIVE 2003, Charlottesville, VA, USA, September 27–October 1, 2003, Revised Selected and Invited Papers, Lecture Notes in Computer Science*, vol. 3062, pp. 446–453. Springer (2003). [https://doi.org/10.1007/978-3-540-25959-6\\_35](https://doi.org/10.1007/978-3-540-25959-6_35)
  54. Taentzer, G., Ohrndorf, M., Lamo, Y., Rutle, A.: Change-preserving model repair. In: *Fundamental Approaches to Software Engineering—20th International Conference, FASE 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22–29, 2017, Proceedings, Lecture Notes in Computer Science*, vol. 10202, pp. 283–299. Springer (2017). [https://doi.org/10.1007/978-3-662-54494-5\\_16](https://doi.org/10.1007/978-3-662-54494-5_16)
  55. Varró, G., Deckwerth, F.: A rete network construction algorithm for incremental pattern matching. In: Duddy, K., Kappel, G. (eds.) *Theory and Practice of Model Transformations*, pp. 125–140. Springer, Berlin (2013)
  56. Wang, M., Gibbons, J., Wu, N.: Incremental updates for efficient bidirectional transformations. In: *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming, ICFP '11*, pp. 392–403. ACM, New York, NY, USA (2011). <https://doi.org/10.1145/2034773.2034825>

57. Weidmann, N., Anjorin, A., Fritsche, L., Varró, G., Schürr, A., Leblebici, E.: Incremental Bidirectional Model Transformation with eMoflon::IBeX. In: Cheney, J., Ko, H. (eds.) Proceedings of the 8th International Workshop on Bidirectional Transformations (BX), CEUR Workshop Proceedings, vol. 2355, pp. 45–55. CEUR-WS.org. <http://ceur-ws.org/Vol-2355/paper4.pdf> (2019)
58. Weidmann, N., Anjorin, A., Robrecht, P., Varró, G.: Incremental (Unidirectional) Model Transformation with eMoflon::IBeX. In: Guerra, E., Orejas, F. (eds.) Graph Transformation—12th International Conference, ICGT 2019, Held as Part of STAF 2019, Eindhoven, The Netherlands, July 15–16, 2019, Proceedings, Lecture Notes in Computer Science, vol. 11629, pp. 131–140. Springer (2019). [https://doi.org/10.1007/978-3-030-23611-3\\_8](https://doi.org/10.1007/978-3-030-23611-3_8)
59. Weidmann, N., Oppermann, R., Robrecht, P.: A feature-based classification of triple graph grammar variants. In: Nierstrasz, O., Gray, J., Oliveira, B.C.d.S. (eds.) Proceedings of the 12th ACM SIGPLAN International Conference on Software Language Engineering, SLE 2019, Athens, Greece, October 20–22, 2019, pp. 1–14. ACM (2019). <https://doi.org/10.1145/3357766.3359529>
60. Xiong, Y., Song, H., Hu, Z., Takeichi, M.: Synchronizing concurrent model updates based on bidirectional transformation. *Softw. Syst. Model.* **12**(1), 89–104 (2013). <https://doi.org/10.1007/s10270-010-0187-3>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.