

# Bidirectional Model Transformation with Precedence Triple Graph Grammars

Technical Report submitted by  
Marius Lauder<sup>\*</sup>, Anthony Anjorin<sup>\*</sup>, Gergely Varró<sup>\*\*</sup>, and Andy Schürr



TECHNISCHE  
UNIVERSITÄT  
DARMSTADT



Real-Time Systems Lab

Department of Electrical Engineering  
and Information Technology (FB18)

Adjunct Member Department of  
Computer Science (FB20)  
Merckstr. 25  
64283 Darmstadt

[www.es.tu-darmstadt.de](http://www.es.tu-darmstadt.de)

<sup>\*</sup> Supported by the ‘Excellence Initiative’ of the German Federal and State Governments and the Graduate School of Computational Engineering at TU Darmstadt.

<sup>\*\*</sup> Supported by the Postdoctoral Fellowship of the Alexander von Humboldt Foundation and associated with the Center for Advanced Security Research Darmstadt.

email: [name.surname@es.tu-darmstadt.de](mailto:name.surname@es.tu-darmstadt.de)

---

Please cite this document as:

URN: [urn:nbn:de:tuda-tuprints-29162](https://nbn-resolving.org/urn:nbn:de:tuda-tuprints-29162)

URL: <http://tuprints.ulb.tu-darmstadt.de/2916/1/29162.pdf>

This document was provided by tuprints,

TU Darmstadt E-Publishing-Service

<http://tuprints.ulb.tu-darmstadt.de>

[tuprints@ulb.tu-darmstadt.de](mailto:tuprints@ulb.tu-darmstadt.de)



This publication complies to the Creative Commons License:

Attribution – Non-Commercial –No Derivative Works 2.0 Germany

<http://creativecommons.org/licenses/by-nc-nd/2.0/de/>

---

---

## Abstract

---

Triple Graph Grammars (TGGs) are a rule-based technique with a formal background for specifying bidirectional model transformation. In practical scenarios, the unidirectional rules needed for the forward and backward transformations are automatically derived from the TGG rules in the specification, and the overall transformation process is governed by a control algorithm. Current implementations either have a worst case exponential runtime complexity or pose such strong restrictions on the class of supported TGGs that practical real-world applications become infeasible. This paper, therefore, introduces a new class of TGGs together with a control algorithm that drops a number of practice-relevant restrictions for TGG rules and still has a polynomial runtime complexity.

---

---

## 1 Introduction

---

The paradigm of Model-Driven Engineering (MDE) has established itself as a promising means of coping with the increasing complexity of modern software systems and, in this context, *model transformation* plays a central role [CH06, SK03]. As industrial applications require reliability and efficiency, the need for formal frameworks that guarantee useful properties of model transformation arises. This is especially the case for *bidirectional* model transformation, where defining a precise semantics for the automatic manipulation and synchronization of models with a corresponding efficient tool support is quite challenging [CFH<sup>+</sup>09]. Amongst the numerous bidirectional model transformation approaches surveyed in [Ste08], the concept of *Triple Graph Grammars (TGGs)* features not only solid formal foundations [EEE<sup>+</sup>07, KLKS10] but also various tool implementations [GHL10, KRW04, KLKS10].

TGGs [Sch94] provide a declarative, rule-based means of (i) specifying the consistency of source and target models in their respective domains, and (ii) tracking inter-domain relationships between model elements explicitly by automatically maintaining a correspondence model. Although TGGs describe how *triples* consisting of source, correspondence, and target models are simultaneously derived, most practical software engineering scenarios require that source or target models already exist and that the models in the correspondence and the opposite domain be consistently constructed by a unidirectional forward or backward transformation. As a consequence, TGG tools that support bidirectional model transformation (i) rely on unidirectional forward and backward operational rules, automatically derived from a single TGG specification, as basic transformation steps, and (ii) use an algorithm that controls which rule is to be applied on which part of the input graph. As a TGG rule in the specification might require *context* elements created by another TGG rule, the control algorithm must consider these dependencies at runtime when (a) determining the order in which graph nodes can be processed, and (b) selecting the rule to be applied.

The unidirectional transformations together with a backtracking control algorithm have already been proven by [EEE<sup>+</sup>07, Sch94] to be (i) *correct* meaning that only consistent graph triples are produced, and (ii) *complete* meaning that all consistent triples, which can be derived from a source or a target graph, can actually be produced. The backtracking algorithm, which uses no look-ahead information for selecting processable nodes and applicable rules, fails in real-world applications due to its exponential runtime. Hence, all TGG approaches employ certain look-ahead strategies to avoid backtracking, trading *expressiveness* (restrictions are posed on the class of supported TGGs) to ensure polynomial runtime.

If supported TGGs are required to exhibit functional behaviour [GHL10, HGO10], which can be enforced statically via a critical pair analysis [EEPT06], the control algorithm cannot make wrong choices enabling an efficient implementation. However, as shown in [KÖ5, RLSS11], transformations required for real-world applications are often non-functional.

To handle a larger set of scenarios including non-functional transformations, a context-driven algorithm has been proposed in [KLKS10], which has been proven to

---

be correct, complete, and polynomial. This algorithm uses an arbitrary node in the existing model as a starting point, and calculates a valid rule sequence by exploring *only* the (transitive) dependencies (required context nodes) of the starting node in a bottom-up manner. This means that this technique lacks a global view on the overall dependencies, which appears to be unsuitable for an *incremental* synchronization scenario. Moreover, this approach is unable to handle dependencies caused by context elements in the domain under construction and transforms only a single node in each transformation step.

To provide a global view on the dependencies in the source graph, a node precedence analysis [LK11] has been introduced, but was only used to statically ensure sufficient conditions for the applicability of the transformation, and not to guide the transformation process itself. Furthermore, this approach was restricted to transformation rules with almost fully connected elements.

In this paper, we extend the node precedence analysis of [LK11] to consider *paths* and not only direct edges between nodes in TGG rules and, therefore, remove the restrictions of [KLKS10, LK11] regarding rule complexity. Additionally, we combine the extended node precedence analysis with a rule dependency analysis to further support the control algorithm of [KLKS10] in determining the node processing order and selecting the next applicable rule. This extended approach can now exploit global dependency information, and perform an iterative, top-down resolution which is more expressive (can handle a larger class of TGGs) and fits better into future incremental scenarios. Finally, we prove that the improved control algorithm is still correct, complete, and polynomial.

Section 2 introduces fundamental definitions using our running example while Sect. 3 discusses and compares existing TGG batch algorithms. Sect. 4.2 presents our rule dependency and node precedence analysis, which is used for the improved TGG batch algorithm presented in Sect. 5. Finally, Sect. 6 complements the discussion in Sect. 3 on related TGG approaches, by giving a broader overview of related *bidirectional* approaches. Sect. 7 concludes with a summary and a discussion of areas for future work.

## 2 Fundamentals and Running Example

In this section, all concepts required to formalize and present our contribution are introduced and explained using our running example.


### 2.1 Type Graphs, Typed Graphs and Triples

We introduce the concept of a *graphs*, and formalize *models* as *typed graphs*.

**Definition 1 (Graph and Graph Morphism)** A graph  $G = (V, E, s, t)$  consists of finite sets  $V$  of nodes, and  $E$  of edges, and two functions  $s, t : E \rightarrow V$  that assign each edge source and target nodes. A graph morphism  $h : G \rightarrow G'$ , with  $G' = (V', E', s', t')$ , is a pair of functions  $h := (h_V, h_E)$  where  $h_V : V \rightarrow V'$ ,  $h_E : E \rightarrow E'$  and  $\forall e \in E : h_V(s(e)) = s'(h_E(e)) \wedge h_V(t(e)) = t'(h_E(e))$ .

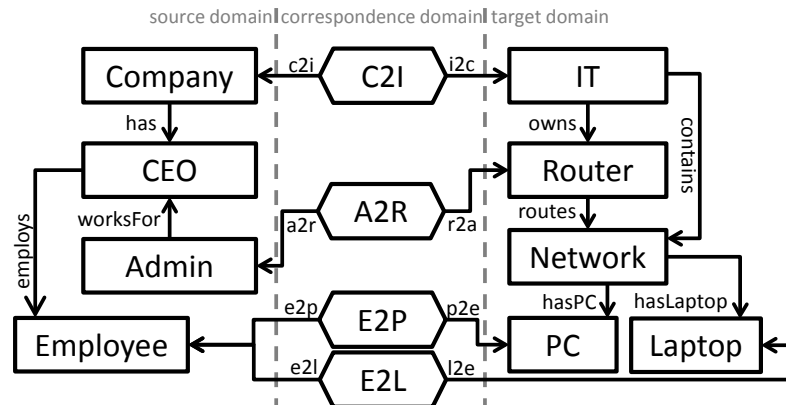
#### Definition 2 (Typed Graph and Typed Graph Morphisms)

A type graph is a graph  $TG = (V_{TG}, E_{TG}, s_{TG}, t_{TG})$ . A pair  $(G, type)$  of a graph  $G$  together with a graph morphism  $type : G \rightarrow TG$ , is called a typed graph. Given typed graphs  $(G, type)$  and  $(G', type')$ ,  $g : G \rightarrow G'$  is a typed graph morphism iff the diagram to the right commutes.



These concepts can be lifted in a straightforward manner to *triples* of connected graphs denoted as  $G = G_S \leftarrow G_C \rightarrow G_T$  as shown by [EEPT06, KLKS10]. In the following, we work with *typed graph triples* and corresponding morphisms.

**Example:** Our running example specifies the integration of *company structures* and corresponding *IT structures*. The *TGG schema* depicted in Fig. 1 is the type graph triple for our running example.



**Figure 1:** TGG Schema for the integration of a company with its IT structure

The *source domain* is described by a type graph for company structures: A Company consists of a CEO, Employees and Admins. In the *target domain*, an IT structure (IT) provides PCs and Laptops in Networks controlled by a Router. The *correspondence domain* specifies valid links between elements in the different domains, in this case a link

between a Company and an IT structure, Admins and Routers, and between Employees and PCs/Laptops.

A schema conform (typed graph) triple is depicted in Fig. 2. The company ES has a CEO named Andy for whom the administrator Ingo works. Additionally, Andy employs Tony and Marius. The corresponding IT structure ES-IT consists of a router WP53 for the network ES-LAN with a PC PC65 and a laptop X200.

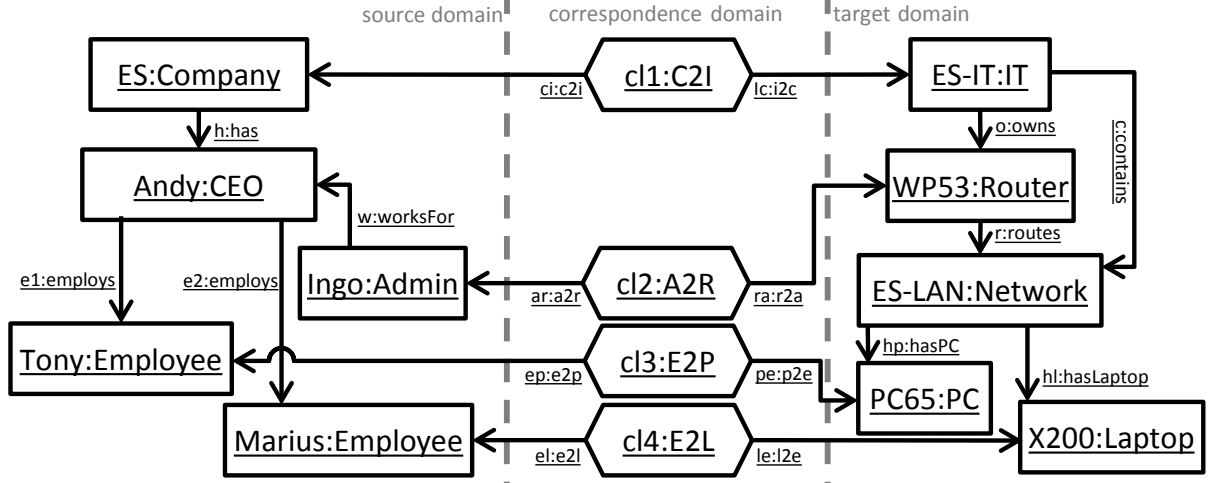



Figure 2: A TGG schema conform triple

## 2.2 Triple Graph Grammars and Rules

The simultaneous evolution of typed graph triples such as our example triple (Fig. 2) can be described by a *triple graph grammar* consisting of *transformation rules*. This is formalized in the following definitions.

### Definition 3 (Graph Triple Rewriting for Monotonic TGG Rules)

A monotonic TGG rule  $r := (L, R)$ , is a pair of typed graph triples such that  $L \subseteq R$ . A TGG rule  $r$  rewrites a graph triple  $G$  into a graph triple  $G'$  via a match  $m : L \rightarrow G$ , denoted as  $G \xrightarrow{r@m} G'$ , iff  $m' : R \rightarrow G'$  is defined by building a pushout as denoted in the diagram to the right.

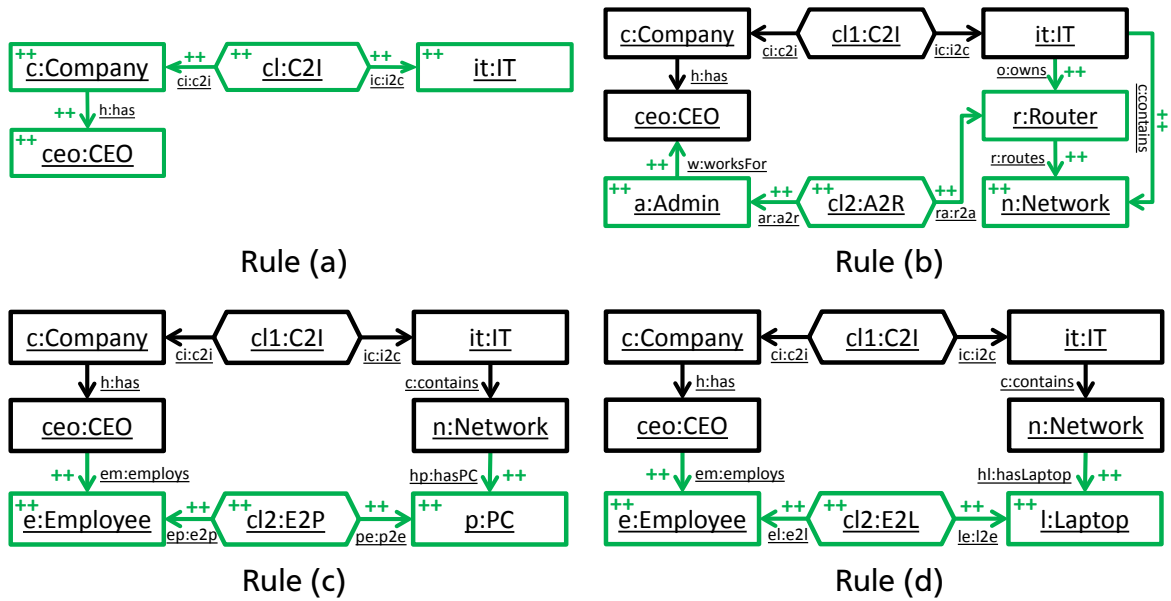


Note that elements in  $L$  denote the precondition of a rule and are referred to as *context elements*, while elements in  $R \setminus L$  are created by a rule and are, therefore, referred to as *created elements*.

**Definition 4 (Triple Graph Grammar)** A triple graph grammar  $TGG := (TG, \mathcal{R})$  is a pair consisting of a type graph triple  $TG$  and a finite set  $\mathcal{R}$  of TGG rules. The generated language is

$\mathcal{L}(TGG) := \{G \mid \exists r_1, r_2, \dots, r_n \in \mathcal{R} : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} \dots \xrightarrow{r_n} G_n = G\}$ ,  
where  $G_0$  denotes the empty graph triple.

**Example:** The four rules depicted in Fig. 3 build up an integrated company and IT structure simultaneously. Rule (a) creates the root elements of the models (a Company with a CEO and a corresponding IT), while Rule (b) appends additional elements (an Admin and a corresponding Router with the controlled Network). Rules (c) and (d) extend the connected models with a new Employee, who can choose a PC or a Laptop. Note that the TGG rules are depicted by merging  $L$  and  $R$  of the rule in a single diagram. Context elements do not have any markup and are black, while created elements are indicated by a “++” markup and are additionally green.



**Figure 3:** Rules (a)–(d) for the integration


### 2.3 Derived Operational Rules

The real potential of TGGs as a bidirectional transformation language lies in the automatic derivation of *operational rules*. Such operational rules can be used to transform a given source domain model to produce a corresponding target domain model and vice versa. Although we focus in the following sections only on a forward transformation, all concepts and arguments are symmetric and can be applied analogously for the case of a backward transformation.

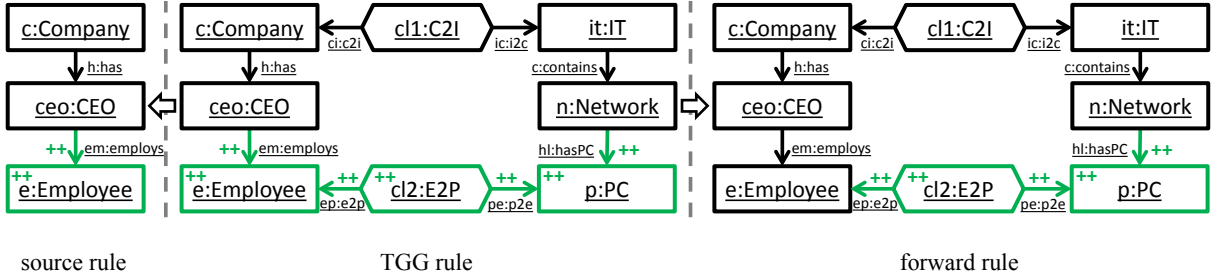
It has been proven by [EEE<sup>+</sup>07, Sch94] that a sequence of TGG rules, which describes a simultaneous evolution, can be uniquely decomposed into (and conversely composed from) a sequence of *source rules* that only evolve the source model and *forward rules* that retain the source model and evolve the correspondence and target models. These operational rules serve as the building blocks used by a control algorithm for unidirectional forward and backward transformation.



**Definition 5 (Derived Operational Rules)** Given a TGG = (TG, R) and a TGG rule  $r = (L, R) \in \mathcal{R}$ , a source rule  $r_S = (SL, SR)$  and a forward rule  $r_F = (FL, FR)$  can be derived according to the following diagram:



**Example:** From Rule (c) of our running example (Fig. 3), the operational rules  $r_S$  and  $r_F$  depicted in Fig. 4 can be derived. The source rule extends the source graph by adding an Employee to an existing CEO in a Company, while the forward rule  $r_F$  transforms an existing Employee of a CEO by creating a new E2P link and PC in the corresponding Network.



**Figure 4:** Source and forward rules derived from Rule (c)

### 3 Related Work on TGG Control Algorithms

Constructing forward (and conversely backward) transformations from operational rules requires a *control algorithm* that is able to determine the correct sequence of forward rules to be applied to a given source graph. The challenge here is to specify a control algorithm that is correct, complete, polynomial, and still expressive enough for real-world applications. To better understand this challenge, we discuss how existing algorithms handle the source graph of our example triple (Fig. 2).

**(I) Bottom-Up, Context-Driven and Recursive:** An established strategy is to transform elements in a bottom-up context-driven manner, i.e., to start with a random node and check if all context nodes are already transformed *before* the selected initial node can be transformed. If a context node is not yet transformed, the algorithm transforms it, recursively checking and transforming its context. Context-driven algorithms always start their transformation process with an arbitrarily selected node, without “knowing” if this was a good choice, i.e., if the node can be transformed immediately or if the input model as a whole is even valid. Such algorithms are correct, but, in general, have problems with completeness due to wrong *local* decisions.

---

**(I.a) Backtracking:** A simple backtracking strategy could be employed to cope with wrong local decisions. For our example, a first iteration over all nodes would determine that only ES together with Andy can be transformed by applying Rule (a). In a second iteration the algorithm would determine again in a trial and error manner that only Ingo can be transformed next with Rule (b), as neither Tony nor Marius can be transformed using Rule (c) or (d) (a Network is missing in the opposite domain). Finally, Tony and Marius can be transformed with either Rule (c) or (d). This algorithm is correct and complete as shown in [EEE<sup>+</sup>07, Sch94] but has exponential runtime and is, therefore, impractical for real-world applications.

It is, however, possible to guarantee polynomial runtime of the context-driven recursion strategy by restricting the class of supported TGGs appropriately as in the following approaches.

**(I.b) Functional Behavior:** Demanding *functional behavior* [GHL10, HGO10] guarantees that the algorithm can choose freely between applicable rules at every decision point and will always get the *same result* without backtracking. Although functional behavior might be suitable for fully automatic integrations, our experience with industrial partners [RLSS11] shows that user interaction or similar guidance (e.g., configuration files) of the integration process is required and leads naturally to non-functional sets of rules with certain degrees of freedom. Please note that our running example is clearly non-functional due to Rules (c) and (d), which can be applied to the same elements on the source side, but create different elements on the target side. Therefore, depending on the choice of rule applications, *different* target graphs are possible with our running example. Demanding functional behavior is a strong restriction that reduces the expressiveness and suitability of TGGs for real-world applications [KLKS10, SK08]. Nevertheless, such a strategy has polynomial runtime and its applicability can be enforced statically.

**(I.c) Local Completeness:** Algorithms that allow a non-functional set of rules still need to guarantee completeness for a certain class of TGGs and, hence, [KLKS10] demands *local completeness*, i.e., that a local decision between rules that can transform the current node *cannot* lead to a dead-end. This means that a local choice (which can be influenced by the user or some other means) might actually result in *different* output graphs, which are, however, always consistent, i.e., in the defined language of the TGG ( $\mathcal{L}(TGG)$ ). For our running example, we could start with an arbitrary node, e.g., Ingo. According to Rule (b), a CEO and a Company are required as context and Rule (a) will thus be applied to ES and Andy. After transforming Ingo, Tony and Marius can be transformed in an arbitrary order, each time making a local choice if a PC (Rule (c)) or Laptop (Rule (d)) is to be created. Furthermore, a *dangling edge check* is introduced in [KLKS10] to further enlarge the class of supported TGGs via a look-ahead to prevent wrong local decisions that would lead to “dangling” edges that can no longer be transformed. Note that our running example is *not* local complete, as it cannot be decided whether an Admin or an Employee should be transformed first (Rules (c) and (d) demand an element on the *target* side that can only be created by Rule (b)). For this reason, the algorithm might fail if it decides to start with one of the Employees. In this case, Rules (c) and (d) would state that ES and Andy are required as context

and have to be transformed first. This is, however, insufficient as a *Network* must be present in the target domain as well. The simple context-driven approach fails here as transforming *ES* and *Andy* with Rule (a) *does not* guarantee that the employees *Marius* and *Tony* can be transformed. The problem here is that context-driven algorithms only regard the given input graph for controlling the rule application and do not consider *cross-domain context dependencies* such as *Network*.

**(II) Top-Down and Iterative:** In contrast to context-driven recursive strategies, algorithms can operate in a top-down iterative manner exploiting a certain *global view* on the whole input graph instead of arbitrarily choosing a node to be transformed.

**(II.a) Correspondence-Driven:** The algorithm presented by [KRW04] requires that all TGG rules demand and create at least one correspondence link, i.e., a hierarchy of correspondence links must be built up during the transformation. The correspondence model can in this case be used to store dependencies between links and is interpreted as a directed acyclic graph, which is used to drive and control the transformation. This algorithm is both batch *and* incremental but it is unclear from [KRW04] for which class of TGGs completeness can be ensured.

**(II.b) Precedence-Driven:** A precedence-driven strategy defines and uses a partial order of nodes in the source graph according to their *precedence*, i.e., the sorting guarantees that the nodes can be transformed in a sequence compatible with the partial order.

---

## 4 Rule Dependency and Precedence Analysis for TGGs

---

In this section, we present a node precedence analysis that provides a partial order required for a precedence-driven strategy, together with a rule dependency analysis that solves the problem of cross-domain context dependencies (as described previously) in some cases.

---


### 4.1 Rule Dependency Analysis

---

To handle cross-domain context dependencies, we utilize the concept of *sequential independence* as introduced by [EEPT06], to statically determine which rules depend on other rules. The intuition is that a rule  $r_2$  depends on another rule  $r_1$  if  $r_1$  creates elements that  $r_2$  requires as context.

#### Definition 6 (Rule Dependency Relation $\prec_R$ )

Given rules  $r_1 = (L_1, R_1)$  and  $r_2 = (L_2, R_2)$ ,  $r_2$  is sequentially dependent on  $r_1$  iff a graph  $D$  and morphisms  $f, h$  exist, such that there exists no morphism  $g$  as depicted to the right, i.e., at least one element required by  $r_2$  (an element in  $L_2$ ), is created by  $r_1$  (this element is in  $R_1$  but not in  $L_1$ ).



The precedence relation  $\prec_R \subseteq \mathcal{R} \times \mathcal{R}$  is defined for a given TGG as follows:

$$r_1 \prec_R r_2 \iff r_2 \text{ is sequentially dependent on } r_1.$$

In practice,  $\prec_R$  can be calculated statically by determining all possible intersections of  $R_1$  and  $L_2$ . If at least one element in an intersection is not in  $L_1$  then  $r_2$  is sequentially dependent on  $r_1$  (i.e.,  $r_1 \prec_R r_2$ ).

**Example:** For the TGG rules of our running example (Fig. 3), the following pairs of rules constitute  $\prec_R$ :  $\text{Rule (a)} \prec_R \text{Rule (b)}$ ,  $\text{Rule (a)} \prec_R \text{Rule (c)}$ ,  $\text{Rule (a)} \prec_R \text{Rule (d)}$ ,  $\text{Rule (b)} \prec_R \text{Rule (c)}$ , and  $\text{Rule (b)} \prec_R \text{Rule (d)}$ .

## 4.2 Precedence Analysis

The following definitions present our path-based node precedence analysis which is used to topologically sort the nodes in a source graph and thus control the transformation process iteratively:


**Definition 7 (Paths and Type Paths)** Let  $G$  be a typed graph with type graph  $TG$ . A path  $p$  between two nodes  $n_1, n_k \in V_G$  is an alternating sequence of nodes and edges in  $V_G$  and  $E_G$ , respectively, denoted as  $p := n_1 \cdot e_1^{\alpha_1} \cdot n_2 \cdot \dots \cdot n_{k-1} \cdot e_{k-1}^{\alpha_{k-1}} \cdot n_k$ , where  $\alpha_i \in \{+, -\}$  specifies if an edge  $e_i$  is traversed from source ( $s(e_i)$ ) to target ( $t(e_i)$ ) (+), or in a reverse direction (-). A type path is a path between node types and edge types in  $V_{TG}$  and  $E_{TG}$ , respectively. Given a path  $p$ , its type is defined as the type path  $\text{type}_p(p) := \text{type}_V(n_1) \cdot \text{type}_E(e_1)^{\alpha_1} \cdot \text{type}_V(n_2) \cdot \text{type}_E(e_2)^{\alpha_2} \cdot \dots \cdot \text{type}_V(n_{k-1}) \cdot \text{type}_E(e_{k-1})^{\alpha_{k-1}} \cdot \text{type}_V(n_k)$ .

For our analysis we are only interested in (type) paths that are induced by certain node creation patterns present in the TGG rules.

**Definition 8 (Relevant Node Creation Patterns)** For a  $TGG = (TG, \mathcal{R})$  the following node creation patterns for paths are defined:


**context(p)**  $\Leftrightarrow \exists r = (L, R) \in \mathcal{R} : p$  is a path between two nodes  $n, n' \in R :$

$(n \in L) \wedge (n' \in R \setminus L)$ , i.e., a rule  $r$  in  $\mathcal{R}$  contains a path  $p$  which is isomorphic to the node creation pattern depicted in the diagram to the right.



**create(p)**  $\Leftrightarrow \exists r = (L, R) \in \mathcal{R} : p$  is a path between two nodes  $n, n' \in R :$

$(n \in R \setminus L) \wedge (n' \in R \setminus L)$ , i.e., a rule in  $\mathcal{R}$  contains a path  $p$  which is isomorphic to the node creation pattern to the right.



We can now define the set of interesting paths, relevant for our analysis.

**Definition 9 (Rule (Type) Path Sets)** For a  $TGG = (TG, \mathcal{R})$ , the rule type path set for the source domain is defined as  $TP_S := \{\text{type}_p(p) \mid \text{context}(p) \vee \text{create}(p)\}$ . For a given typed graph  $G$ , the rule path set for the source domain is defined as  $P_S := \{p \mid p \text{ is a path between } n, n' \in V_G \wedge \text{type}_p(p) \in TP_S\}$ .

The following definitions formalize the concept of *precedence* between nodes.

**Definition 10 (Precedence Function  $\mathcal{PF}_S$ )** Let  $\mathcal{P} := \{<, \dot{=}, \cdot\}$  be the set of precedence relation symbols. Given a TGG  $= (TG, \mathcal{R})$  and the rule type path set for the source domain  $TP_S$ . The precedence function for the source domain  $\mathcal{PF}_S : TP_S \rightarrow \mathcal{P}$  is computed as follows:

$$\begin{aligned} &< \text{ iff } \text{context}(p) \wedge \neg \text{create}(p) \\ \mathcal{PF}_S(\text{type}_p(p)) &:= \dot{=} \text{ iff } \neg \text{context}(p) \wedge \text{create}(p) \\ &\cdot \text{ otherwise} \end{aligned}$$

**Example:**  $\mathcal{PF}_S$  for our running example consists of the following entries:

Rule (a):  $\mathcal{PF}_S(\text{Company} \cdot \text{has}^+ \cdot \text{CEO}) = \dot{=}$

Rule (b):  $\mathcal{PF}_S(\text{Company} \cdot \text{has}^+ \cdot \text{CEO} \cdot \text{worksFor}^- \cdot \text{Admin}) = <$  and

$$\mathcal{PF}_S(\text{CEO} \cdot \text{worksFor}^- \cdot \text{Admin}) = <$$

Rules (c) and (d):  $\mathcal{PF}_S(\text{Company} \cdot \text{has}^+ \cdot \text{CEO} \cdot \text{employs}^- \cdot \text{Employee}) = <$  and

$$\mathcal{PF}_S(\text{CEO} \cdot \text{employs}^- \cdot \text{Employee}) = <$$

Based on the precedence function  $\mathcal{PF}_S$ , relations  $<_S$  and  $\dot{=}_S^*$  can now be defined and used to topologically sort a given input graph and determine the sets of elements that can be transformed at each step in the algorithm.

**Definition 11 (Precedence Relation  $<_S$ )** Given  $\mathcal{PF}_S$ , the precedence function for a given TGG, and a typed source graph  $G_S$ . The precedence relation  $<_S \subseteq V_{G_S} \times V_{G_S}$  for the source domain is defined as follows:  $n <_S n'$  if there exists a path  $p \in P_S$  between nodes  $n$  and  $n'$  such that  $\mathcal{PF}(\text{type}_p(p)) = <$ .

**Example:** For our example triple (Fig. 2), the following pairs constitute  $<_S$ :  $(\text{ES} <_S \text{Ingo}), (\text{ES} <_S \text{Tony}), (\text{ES} <_S \text{Marius}), (\text{Andy} <_S \text{Ingo}), (\text{Andy} <_S \text{Tony})$ , and  $(\text{Andy} <_S \text{Marius})$ .

**Definition 12 (Relation  $\dot{=}_S$ )** Given  $\mathcal{PF}_S$ , the precedence function for a given TGG, and a typed source graph  $G_S$ . The symmetric relation  $\dot{=}_S \subseteq V_{G_S} \times V_{G_S}$  for the source domain is defined as follows:  $n \dot{=}_S n'$  if there exists a path  $p \in P_S$  between nodes  $n$  and  $n'$  such that  $\mathcal{PF}(\text{type}_p(p)) = \dot{=}$ .

**Definition 13 (Equivalence Relation  $\dot{=}_S^*$ )** The equivalence relation  $\dot{=}_S^*$  is the transitive and reflexive closure of  $\dot{=}_S$ , which is already symmetric by definition.

**Example:** For our example triple (Fig. 2), the following equivalence classes constitute  $\dot{=}_S^*$ :  $\{\text{Andy}, \text{ES}\}, \{\text{Ingo}\}, \{\text{Tony}\}$ , and  $\{\text{Marius}\}$ .

**Definition 14 (Precedence Graph  $\mathcal{PG}_S$ )** The precedence graph for a given source graph  $G_S$  is a graph  $\mathcal{PG}_S$  constructed as follows:

(i) The equivalence relation  $\dot{=}_S^*$  is used to partition  $V_{G_S}$  into equivalence classes

$$A_1, \dots, A_n \text{ which serve as the nodes of } \mathcal{PG}_S, \text{ i.e., } V_{\mathcal{PG}_S} := \{A_1, \dots, A_n\}.$$

(ii) The edges in  $\mathcal{PG}_S$  are defined as follows:

$$E_{\mathcal{PG}_S} := \{e \mid s(e) = A_i, t(e) = A_j : \exists n_i \in A_i, n_j \in A_j \text{ with } n_i <_S n_j\}.$$

**Example:** The corresponding  $\mathcal{PG}_S$  constructed from our example triple is depicted in Fig. 5(a) in Sect. 5.

---

## 5 Precedence TGG Batch Algorithm

---

In this section, we present our batch algorithm and explain how the introduced rule dependency and node precedence analysis are used to efficiently transform a given source graph. For a forward transformation (a backward transformation works analogously), the input for the algorithm is a graph  $G_S$ , the statically derived rule dependency relation  $\prec_R$ , and the precedence function for the source domain  $\mathcal{PF}_S$ .

---

### Algorithm 1 Precedence TGG Batch Algorithm

---

```
1: procedure TRANSFORM( $G_S, \prec_R, \mathcal{PF}_S$ )
2:    $\mathcal{PG}_S \leftarrow$  BUILDPRECEDENCEGRAPH( $G_S, \mathcal{PF}_S$ )
3:   while ( $\mathcal{PG}_S$  contains equivalence classes) do
4:      $readyNodes \leftarrow$  all nodes in equiv. classes in  $\mathcal{PG}_S$  without incoming edges
5:      $readyNodes \leftarrow$  sort  $readyNodes$  utilizing  $\prec_R$ 
6:     for (node  $n$  in  $readyNodes$ ) do
7:        $transformedNodes \leftarrow$  CHOOSEANDAPPLYRULE( $n$ )
8:       if  $transformedNodes \neq \emptyset$  then
9:          $\mathcal{PG}_S \leftarrow$  remove all nodes in  $transformedNodes$  from  $\mathcal{PG}_S$ 
10:        break
11:       end if
12:     end for
13:     if  $transformedNodes = \emptyset$  then
14:       terminate with error  $\triangleright$  Local Completeness Criterion violated
15:     end if
16:   end while
17:   return  $G_S \leftarrow G_C \rightarrow G_T$ 
18: end procedure
```

---

TRANSFORM (Algorithm 1) determines a graph triple  $G_S \leftarrow G_C \rightarrow G_T$  as output. The first step (line (2)) of the algorithm is to build the precedence graph  $\mathcal{PG}_S$  according to Def. 14. Note that the procedure BUILDPRECEDENCEGRAPH will terminate with an error if there is a cycle in the precedence graph and it is thus impossible to sort the elements of the source graph according to their dependencies. Starting on line (3), a while-loop iterates over equivalence classes in  $\mathcal{PG}_S$  until there are none left. In the while-loop, the set  $readyNodes$  contains all nodes that can be transformed next, i.e., whose context elements have already been transformed (line (4)). This set is determined by taking all nodes in the equivalence classes of  $\mathcal{PG}_S$ , which do not have incoming edges (dependencies). On line (5),  $readyNodes$  is sorted according to the partially ordered relation  $\prec_R$ , i.e., the rules that can be used to transform nodes in  $readyNodes$  are determined, sorted with  $\prec_R$  and reflected in  $readyNodes$ . This could be achieved by assigning an integer to each rule according to the partial order of  $\prec_R$  and then selecting the largest number of all rules that translate  $n \in readyNodes$  for  $n$ .<sup>1</sup> Next, a for-loop iterates over the sorted

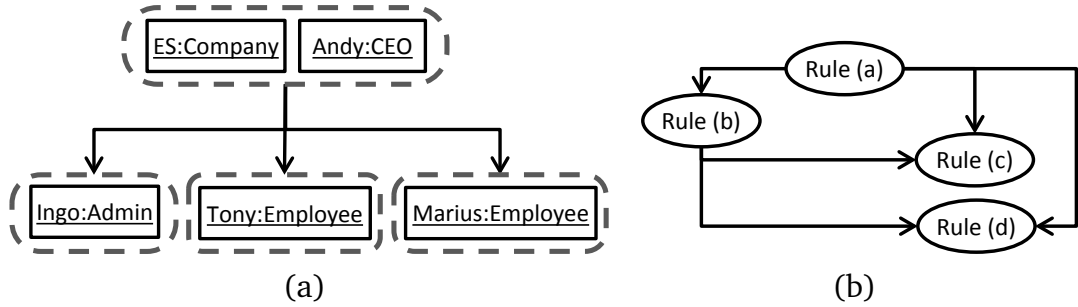
---

<sup>1</sup> If it is not possible to sort  $readyNodes$  due to cycles in  $\prec_R$ , this additional analysis supplies no further information and  $readyNodes$  remains unchanged.



*readyNodes* (line (6)). On line (7) the procedure `CHOOSEANDAPPLYRULE` is used to determine and filter the rules as presented in [KLKS10], allowing for user input or choosing arbitrarily from the final applicable rules. If a rule could be successfully chosen and applied to transform  $n$  on line (7), a non-empty set of *transformedNodes* is returned that is used to update  $\mathcal{PG}_S$  on line (9). In this case, the for-loop is terminated and the while-loop is repeated with the updated and thus “smaller”  $\mathcal{PG}_S$ . If *transformedNodes* is empty, the for-loop is repeated for the next node in *readyNodes*. If *transformedNodes*, however, *remains* empty on line (13), we know that no node in *readyNodes* has been transformed and that the algorithm has hit a dead-end. This can only happen for TGGs that violate the *Local Completeness Criterion* (cf. algorithm strategy I.c in Sect. 3) and are *not* in the class of supported TGGs.

**Example:** To demonstrate the presented algorithm, we apply a forward transformation for the source graph of our example triple depicted in Fig. 2. Given as input is  $G_S$ , the rule dependency relation  $\prec_R$  (depicted as a graph in Fig. 5(b)), and the precedence function  $\mathcal{PF}_S$  (cf. example for Def. 10). On line (2), the precedence graph  $\mathcal{PG}_S$  for  $G_S$ , depicted in Fig. 5(a), is built.  $\mathcal{PG}_S$  is acyclic, hence the transformation can continue.



**Figure 5:**  $\mathcal{PG}_S$  for the input graph (left) and relation  $\prec_R$  for all rules (a)–(d) (right)

On line (4), the set *readyNodes* is determined, consisting in this case of the nodes ES and Andy from a single equivalence class of  $\mathcal{PG}_S$ . On line (5), only one rule can be used to transform both nodes and, therefore, the sorting is trivial. On line (6) ES or Andy is chosen randomly, and in either case, the only candidate rule is Rule (a) (Fig. 3), which can be directly applied on line (7). Again in either case, *transformedNodes* contains both nodes as Rule (a) transforms ES and Andy simultaneously.  $\mathcal{PG}_S$  is updated on line (9) to consist of three unconnected equivalence classes Ingo, Tony, and Marius, and the for-loop terminates. In the second iteration through the while-loop, *readyNodes* now contains all these three elements and will be sorted according to  $\prec_R$  on line (5). This time, the sorting reveals that Ingo must be transformed before Tony and Marius as Rules (c) and (d) both require a Network as context in the target domain, which can only be created by applying Rule (b) first, i.e.,  $Rule (b) \prec_R Rule (c)$ ,  $Rule (b) \prec_R Rule (d)$  (Fig. 5(b)). The for-loop in line (6), therefore, starts with Ingo. Applying Rule (b) (line (7)) puts Ingo in *transformedNodes*,  $\mathcal{PG}_S$  is updated on line (9) to now contain only Tony and Marius and the for-loop is terminated with the break on line (10). In the third iteration, *readyNodes* contains Tony and Marius, and no sorting is needed as Rules (c) and (d) do not depend on each other. Line (6) could randomly select Tony

first and (arbitrarily or via user input) Rule (c) could be chosen to be applied in line (7). After updating  $\mathcal{PG}_S$  again and breaking out of the for-loop, only Marius remains untransformed. Similar to the penultimate iteration, Rule (d) could be selected and applied this time. Updating  $\mathcal{PG}_S$  on line (9) empties the precedence graph, which terminates the while-loop on line (3). The created graph triple depicted in Fig. 2 is returned on line (17).

---

### Formal Properties of the Precedence TGG Batch Algorithm

---

In the following we argue that the presented algorithm retains all formal properties stipulated in [SK08] and proved for the context-driven algorithm of [KLKS10].

#### Definition 15 (Correctness, Completeness and Efficiency)

**Correctness:** Given a source graph  $G_S$ , the transformation algorithm either terminates with an error or produces a graph triple  $G_S \leftarrow G_C \rightarrow G_T \in \mathcal{L}(TGG)$ .

**Completeness:** For all triples  $G_S \leftarrow G_C \rightarrow G_T \in \mathcal{L}(TGG)$ , the transformation algorithm produces a consistent triple  $G_S \leftarrow G'_C \rightarrow G'_T \in \mathcal{L}(TGG)$  for the input source graph  $G_S$ .

**Efficiency:** According to [SK08], a TGG batch transformation algorithm is efficient if its runtime complexity class is  $O(n^k)$ , where  $n$  is the number of nodes in the source graph to be transformed and  $k$  is the largest number of elements to be matched by any rule  $r$  of the given TGG.

All properties are defined analogously for backward transformations.

**Theorem** Algorithm 1 is correct, complete and efficient.

#### Proof

**Correctness:** If the algorithm returns a graph triple without terminating with an error, then it was able to determine a sequence of source rules  $r_{1_S}, r_{2_S}, \dots, r_{n_S}$  that would build the given source graph  $G_S$  and, thus, the corresponding sequence of forward rules  $r_{1_F}, r_{2_F}, \dots, r_{n_F}$  that transform the given source graph (Def. 5). The *Decomposition and Composition Theorem* of [EEE<sup>+</sup>07] guarantees that it is possible to compose the sequence  $r_{1_S}, r_{2_S}, \dots, r_{n_S}, r_{1_F}, r_{2_F}, \dots, r_{n_F}$  to the sequence of TGG rules  $r_1, r_2, \dots, r_n$  which proves that the resulting graph triple is consistent, i.e.,  $G_S \leftarrow G_C \rightarrow G_T \in \mathcal{L}(TGG)$ .  $\square$

**Completeness:** Showing completeness is done in two steps:

First of all, we consider the algorithm without the additional concept of rule dependencies via the relation  $\prec_R$ . The remaining algorithm transforms nodes with the same concepts (e.g., dangling edge check) as the previous algorithm in [KLKS10], but iteratively in a *fixed sequence*, for which we guarantee, by definition of the precedence graph (cf. 14), that the context of every node is always transformed first. As the context-driven strategy taken by the algorithm in [KLKS10] is able to transform a model by arbitrarily choosing an element and transforming its context elements in a bottom-up manner (cf. Sect. 3), the fixed sequence taken by our algorithm must be a possible sequence that could be chosen by the algorithm in [KLKS10]. Algorithm 1 can, therefore, be seen as forcing the context-driven algorithm to transform elements in one of the possible



sequences, from which it can arbitrarily choose. This shows that all completeness arguments from [KLKS10] can be transferred to the new algorithm, i.e., Algorithm 1 is complete for the class of local complete TGGs.

In a second step, we now consider the algorithm with the additional relation  $\prec_R$  and, therefore, the capability of handling specifications with cross-domain context dependencies as in our running example. We have shown in Sect. 3 that the algorithm presented in [KLKS10] cannot cope with such specifications as they violate the local-completeness criterion. We can, hence, conclude that Algorithm 1 is more expressive than the previous context-driven algorithm as it can handle certain TGGs that are not local complete. We leave the precise categorization of this new class of TGGs to future work.  $\square$

**Efficiency:** Building the precedence graph  $\mathcal{PG}_S$  on line (2), essentially a topological sorting, is realizable in  $O(n^l)$ , where  $l$  is the maximum length of relevant paths according to  $\mathcal{PF}_S$ . Note that  $l$  can be at most of size  $k$  (the largest number of elements to be matched by any rule  $r$  of the given TGG), thus we can estimate this with  $O(n^k)$ . The while-loop starting on line (3) iterates through  $\mathcal{PG}_S$ , which will be decreased every time by at least one node from an equivalence class. The while-loop is, thus, run in the worst-case (equivalence classes in  $\mathcal{PG}_S$  all consist of exactly one node)  $n$  times. In the while-loop, we select equivalence classes without incoming edges in line (4). This can be achieved in  $O(n)$  by iterating through  $\mathcal{PG}_S$ . Building the topological order on line (5) requires inspecting all nodes in *readyNodes* and their appropriate rules in  $O(n)$ . The for-loop starting on line (6) iterates in the worst-case over all nodes in *readyNodes* where updating  $\mathcal{PG}_S$  on line (9), requires traversing all successor nodes which is at most  $n - 1$  (i.e.,  $O(n)$ ). As argued in [KLKS10], transforming a node, i.e., checking all conditions and performing pattern matching (line (7)), is assumed to run in  $O(n^k)$  (cf. Def. 15). Summarizing, we obtain:  $n^k + n \cdot (n + n + n \cdot (n^k + n)) \in O(n^k)$ .  $\square$

---

## 6 Related Work on Alternative Bidirectional Languages

---

Complementing our related work on TGG batch algorithms (cf. Sect. 3), we now focus on *alternative bidirectional languages* that share and address similar challenges as TGGs but take fundamentally different strategies. As bidirectionality is a challenge in various application domains and communities, there exists a substantial number of different approaches, formalizations and tools [Ste08]. The lenses framework is of particular interest when compared to TGGs, as [HEO<sup>+</sup>11] has shown that incremental TGGs can be viewed as an implementation of a *delta-based* framework for *symmetric lenses*. Although we have presented a batch algorithm for TGGs, our ultimate goal is to provide a solid basis for an efficient incremental TGG implementation. As compared to existing lenses implementations for string data or trees such as *Boomerang* [BFP<sup>+</sup>08], TGGs are better suited for MDE where model transformations operate on complex *graph-like* structures. Similar to TGGs, *GRoundTram*, a bidirectional framework based on graph transformations [HHI<sup>+</sup>11], aims to support model transformations in the context of MDE. There are, however, a number of interesting differences: (i) While *GRoundTram*

---

demands a forward transformation from the user and automatically generates a consistent backward transformation, TGGs (in this respect similar to lenses) provide a language from which both forward and backward transformations are automatically derived. Both approaches face a different set of non-trivial challenges. (ii) GRoundTram uses UnQL+, basing on the graph query algebra UnCAL, with a strong emphasis on compositionality, while TGGs are rule-based algebraic graph transformations. (iii) GRoundTram maintains traceability in an implicit manner while TGGs create explicit typed traceability links between integrated models, which can be used to store extra information for incremental model synchronization or manual reviews. In contrast to both Boomerang and GRoundTram, TGGs adhere to the fundamental *unification* principle in MDE (everything is a model) and as such, a bidirectional model transformation specified as a TGG is a model which is conform to a well defined TGG metamodel. Unification has wide-reaching consequences including enabling a natural *bootstrap* and *higher order transformations*. Finally, TGGs served as an inspiration and basis for the standard OMG bidirectional transformation language QVT and can be regarded as a valid implementation thereof [Ste08].

---

## 7 Conclusion and Future Work

---

In this paper, an improvement of our previous TGG batch algorithm was presented. We introduced a rule dependency and node precedence analysis of TGG specifications to enable an iterative batch transformation strategy in a top-down manner with increased expressiveness. We have shown that this algorithm runs in polynomial runtime and complies to the formal properties for TGG implementations according to [SK08]. As a next step, we shall implement the presented algorithm as an extension of our current batch implementation in our metamodeling tool eMoflon<sup>2</sup>[ALPS11], and start working on an efficient incremental TGG algorithm based on our rule dependency and node precedence analysis. Last but not least, providing a *rule checker* that decides at compile time if a given TGG can be transformed by our algorithm is a crucial task to improve the usability of our tool.

---

<sup>2</sup> [www.moflon.org](http://www.moflon.org)

---

## References

---

- [ALPS11] Anthony Anjorin, Marius Lauder, Sven Patzina, and Andy Schürr. eMoflon: Leveraging EMF and Professional CASE Tools. In *3. Workshop Methodische Entwicklung von Modellierungswerkzeugen (MEMWe 2011)*, volume 192 of *Lecture Notes in Informatics (LNI)*, Bonn, 2011. Gesellschaft für Informatik.
- [BFP<sup>+</sup>08] Aaron Bohannon, J. Nathan Foster, Benjamin C. Pierce, Alexandre Pilkiewicz, and Alan Schmitt. Boomerang: Resourceful Lenses for String Data. *ACM SIGPLAN Notices*, 43(1):407–419, 2008.
- [CFH<sup>+</sup>09] Krzysztof Czarnecki, John Nathan Foster, Zhenjiang Hu, Ralf Lämmel, Andy Schürr, and James Terwilliger. Bidirectional Transformations: A Cross-Discipline Perspective. In *2<sup>nd</sup> International Conference on Model Transformation (ICMT 2009)*, volume 5563 of *Lecture Notes in Computer Science (LNCS)*, pages 260–283, Berlin / Heidelberg, 2009. Springer.
- [CH06] Krzysztof Czarnecki and Simon Helsen. Feature-Based Survey of Model Transformation Approaches. *IBM Systems Journal*, 45(3):621–645, 2006.
- [EEE<sup>+</sup>07] Hartmut Ehrig, Karsten Ehrig, Claudia Ermel, Frank Hermann, and Gabriele Taentzer. Information Preserving Bidirectional Model Transformations. In *10<sup>th</sup> International Conference on Fundamental Approaches to Software Engineering (FASE 2007)*, volume 4422 of *Lecture Notes in Computer Science (LNCS)*, pages 72–86. Springer, Berlin / Heidelberg, 2007.
- [EEPT06] Hartmut Ehrig, Karsten Ehrig, Ulrike Prange, and Gabriele Taentzer. *Fundamentals of Algebraic Graph Transformation (Monographs in Theoretical Computer Science. An EATCS Series)*. Springer, Berlin / Heidelberg, 2006.
- [GHL10] Holger Giese, Stephan Hildebrandt, and Leen Lambers. Toward Bridging the Gap Between Formal Semantics and Implementation of Triple Graph Grammars. In *2010 Workshop on Model-Driven Engineering, Verification, and Validation (MoDeVVA 2010)*, pages 19–24. IEEE, 2010.
- [HEO<sup>+</sup>11] Frank Hermann, Hartmut Ehrig, Fernando Orejas, Krzysztof Czarnecki, Zinovy Diskin, and Yingfei Xiong. Correctness of Model Synchronization Based on Triple Graph Grammars. In *14<sup>th</sup> International Conference on Model Driven Engineering Languages and Systems (MODELS 2011)*, volume 6981 of *Lecture Notes in Computer Science (LNCS)*, pages 668–682. Springer, Berlin / Heidelberg, 2011.
- [HGO10] Frank Hermann, Ulrike Golas, and Fernando Orejas. Efficient Analysis and Execution of Correct and Complete Model Transformations Based on Triple Graph Grammars. In *1<sup>st</sup> International Workshop on Model-Driven Interoperability (MDI 2010)*, volume 482 of *International Conference Proceedings Series (ICPS)*, pages 22–31. ACM, 2010.

- 
- [HHI<sup>+</sup>11] Soichiro Hidaka, Zhenjiang Hu, Kazuhiro Inaba, Hiroyuki Kato, and Keisuke Nakano. GRoundTram: An Integrated Framework for Developing Well-Behaved Bidirectional Model Transformations. In *26<sup>th</sup> International Conference on Automated Software Engineering (ASE 2011)*, pages 480–483. IEEE, 2011.
- [KÖ5] Alexander Königs. Model Transformation with Triple Graph Grammars. In *Model Transformations in Practice (MTIP 2005)*, 2005.
- [KLKS10] Felix Klar, Marius Lauder, Alexander Königs, and Andy Schürr. Extended Triple Graph Grammars with Efficient and Compatible Graph Translators. In *Graph Transformations and Model Driven Engineering - Essays Dedicated to Manfred Nagl on the Occasion of his 65th Birthday*, volume 5765 of *Lecture Notes in Computer Science (LNCS)*, pages 141–174. Springer, Berlin / Heidelberg, 2010.
- [KRW04] Ekkart Kindler, Vladimir Rubin, and Robert Wagner. An Adaptable TGG Interpreter for In-Memory Model Transformations. In *2<sup>nd</sup> International Fujaba Days (Fujaba Days 2004)*, pages 35–38, 2004.
- [LK11] Marius Lauder and Felix Klar. Precedence Triple Graph Grammars. In *4<sup>th</sup> International Symposium on Applications of Graph Transformation With Industrial Relevance (AGTIVE 2011)*, Online-Proceedings, pages 297–312, Budapest, 2011.
- [RLSS11] Sebastian Rose, Marius Lauder, Michael Schlereth, and Andy Schürr. A Multidimensional Approach for Concurrent Model Driven Automation Engineering. In *Model-Driven Domain Analysis and Software Development: Architectures and Functions*, pages 90–113, Hershey, 2011. IGI Global.
- [Sch94] Andy Schürr. Specification of Graph Translators with Triple Graph Grammars. In *20<sup>th</sup> International Workshop on Graph-Theoretic Concepts in Computer Science (WG 1994)*, volume 903 of *Lecture Notes in Computer Science (LNCS)*, pages 151–163, Berlin / Heidelberg, 1994. Springer.
- [SK03] Shane Sendall and Wojtek Kozaczynski. Model Transformation: The Heart and Soul of Model-driven Software Development. *IEEE Software*, 20(5):42–45, 2003.
- [SK08] Andy Schürr and Felix Klar. 15 Years of Triple Graph Grammars. In *4<sup>th</sup> International Conference on Graph Transformations (ICGT 2008)*, volume 5214 of *Lecture Notes in Computer Science (LNCS)*, pages 411–425. Springer, Berlin / Heidelberg, 2008.
- [Ste08] Perdita Stevens. A Landscape of Bidirectional Model Transformations. In *2<sup>nd</sup> International Summer School on Generative and Transformational Techniques in Software Engineering (GTTSE 2007)*, volume 5235 of *Lecture Notes in Computer Science (LNCS)*, pages 408–424, Berlin / Heidelberg, 2008. Springer.