

A Solution to the Flowgraphs Case Study using Triple Graph Grammars and eMoflon

Anthony Anjorin*

Technische Universität Darmstadt
Real-Time Systems Lab
Merckstr. 25
64283 Darmstadt, Germany

anjorin@es.tu-darmstadt.de

Marius Lauder*

Technische Universität Darmstadt
Real-Time Systems Lab
Merckstr. 25
64283 Darmstadt, Germany

lauder@es.tu-darmstadt.de

After 20 years of Triple Graph Grammars (TGGs) and numerous actively maintained implementations, there is now a need for challenging examples and success stories to show that TGGs can be used for real-world bidirectional model transformations. Our primary goal in recent years has been to increase the expressiveness of TGGs by providing a set of pragmatic features that allow a controlled fallback to programmed graph transformations and Java.

Based on the *Flowgraphs* case study [7] of the Transformation Tool Contest (TTC 2013), we present (i) *attribute constraints* used to express complex bidirectional attribute manipulation, (ii) *binding expressions* for specifying arbitrary context relationships, and (iii) *post-processing methods* as a black box extension for TGG rules. In each case, we discuss the enabled trade-off between guaranteed formal properties and expressiveness.

Our solution, implemented with eMoflon (www.emoflon.org) our metamodeling and model transformation tool, is available as a virtual machine hosted on Share [15].

1 Introduction and Motivation

Triple Graph Grammars (TGGs) [13] are a declarative, rule-based, bidirectional model transformation language and can be used to specify a *consistency relation* over source and target models, with which various *operational scenarios*, such as a forward and backward transformation, can be automatically supported. As TGGs have been in use for about 20 years and numerous formal results based on algebraic graph transformation [8] have been established, it is now time to tackle convincingly realistic real-world case studies with TGGs.

In our opinion, based on practical experience with our TGG implementation in eMoflon [1], what is necessary is to increase the expressiveness of TGGs by providing a pragmatic set of features that allow a controlled *fallback* to a general purpose (transformation) language such as standard Java or programmed graph transformations via, e.g., *Story Driven Modelling* (SDM) [4].

Our contribution in this paper is to present a set of advanced TGG features that enable a controlled integration of SDM and Java code in TGG rules. These features are: (i) *attribute constraints* for complex bidirectional attribute manipulation in TGG rules, (ii) *binding expressions* for expressing complex, possibly recursive context relationships, and (iii) *post-processing methods* that allow for arbitrary post-processing after a TGG rule has been applied. Using the case-study *Flowgraphs* of the TTC 2013 [7], we present each feature, explaining when it is necessary and discussing the involved trade-off between expressiveness and formal properties.

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

2 Solution With Triple Graph Grammars

The *Flowgraphs* case study of the TTC 2013 is a text-to-model transformation involving the static analysis of a small Java subset. The task is to transform a simple Abstract Syntax Tree (AST) for Java to a *flowgraph*, i.e., a model containing explicit information about the control and data flow in the program. In this context, *bidirectionality* is a crucial requirement to reflect program manipulations (refactorings, quick-fixes), applied via model transformations, back in the code. To show that eMoflon supports using TGGs directly with standard parser and unparser technology (including XML), we modify the case study slightly by using ANTLR [12] and StringTemplate [11] for parsing and unparsing Java code, respectively. Our TGG rules operate directly on the AST produced by ANTLR without a model parser such as EMFText [5].

A TGG consists of a set of *rules* describing how source and target models evolve simultaneously. Each *TGG rule* consists of *elements* (objects and links), typed according to the classes and references in the respective metamodels. The TGG rule MethodRule (Fig. 1) creates the root element of the AST representing a class containing a single method (on the left hand side), and a Method with an explicit Exit element in the model (on the right hand side). As can be seen from the rule, the AST metamodel (not shown explicitly) consists basically of labelled nodes with children and attributes.

MethodRule is referred to as an axiom as it only creates elements (green with a “++” markup) and does not require any context (precondition). Note how *correspondence* elements such as nodeToBlock, depicted visually as hexagons, are created to connect AST elements (classNode) with model elements (method). These elements are used in other rules as traceability links and conform to a simple correspondence metamodel, connecting source and target metamodels. The triple of source, correspondence, and target metamodels is referred to as a *TGG schema*.

Simple attribute manipulation can be expressed as *attribute assignments* such as `txt = "Exit"` in the exit element. Note that this is automatically interpreted as an assertion for the backward transformation, where the model is parsed and the AST is created.

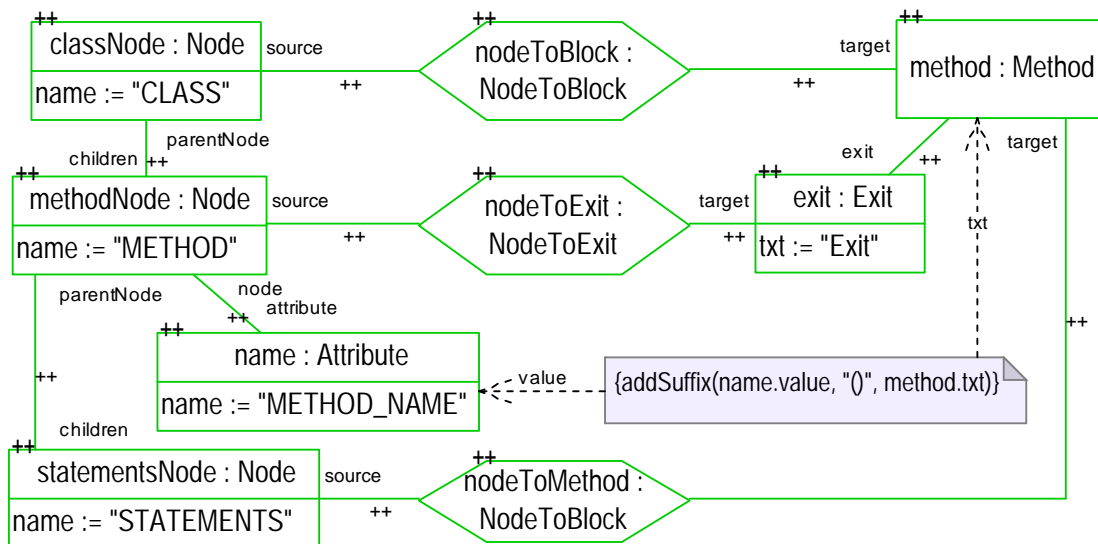


Figure 1: TGG axiom MethodRule for creating the corresponding tree structure for Methods

More complex attribute manipulation can be specified using bidirectional *attribute constraints*, such as the constraint `addSuffix(name.value, "()", method.txt)` in `MethodRule`, which expresses that the value of the name attribute in the AST must be equal to the `txt` attribute of the method element in the model, after adding `()` as a suffix. Attribute constraints such as `addSuffix` are interpreted appropriately when compiling the TGG rule for each operational scenario (e.g., forward, backward) and can thus be expressed in a direction agnostic manner fitting to the rest of the TGG rule.

A second TGG rule `AssignmentWithExpRule` (Fig. 2) handles, for example, statements of the form `"a = b + 3;"` represented in the AST as an assignment tree with `lhs = "a"`, `rhs = "+"`, `operandL = "b"`, and `operandR = "3"`. To test the flexibility of the transformation language, the case study requires an embedded *model-to-text* transformation to transform the tree structure back to text and store it as the attribute value of the `SimpleStmt` element in the model. This is accomplished in `AssignmentWithExpRule` with a set of attribute constraints, referred to as the *attribute Constraint Satisfaction Problem (CSP)* of the TGG rule. How this works is best explained using a concrete example; To handle `"a = b + 3;"` in the forward transformation, the set of constraints is sorted and solved as follows:

```
isAnIdentifier("a") => true
concatWithOperatorSymbol("+", "b", "3", temp1) => temp1 = "b + 3"
concat("=", "a", "b + 3", temp2) => temp2 = "a = b + 3"
addSuffix("a = b + 3", ";", declStm.txt) => declStm.txt = "a = b + 3;"
```

The constraints `isAnIdentifier` and `concatWithOperatorSymbol` are *user defined constraints* and were implemented in Java specifically for the Flowgraphs case study. All other constraints are *library constraints* and are directly available for use in our tool. In this way, the set of constraints can be seamlessly extended by the user, reused, and combined with other constraints in different TGG rules. The constraint `isAnIdentifier` ensures that the rule is not applied for statements such as `"int a = b + 3"`, where `"int a"` is not an identifier and is handled with a different rule for declarations. The second user defined constraint `concatWithOperatorSymbol` is an extension of the normal `concat` constraint and is able to split expressions using a list of supported operators as potential characters for splitting/concatenation as required. The attribute CSP is sorted and solved differently but analogously for the backward transformation.

Formally, attribute CSPs serve as *application conditions* for TGG rules, and, as access is restricted to only attribute values, most formal results for TGGs still hold as we have shown in previous work [2]. Attribute constraints are implemented in Java and thus allow a controlled integration of Java code for complex attribute manipulation in TGG rules.

Note how the rule requires context elements (black without any markup) that must already exist and have been created by applying a different rule such as `MethodRule`. In the model, `AssignmentWithExpRule` requires (i) the parent block element that is to contain the created assignment statement `declStatement`, and (ii) the succeeding element `nextFlow` in the flowgraph.

Although the created `cfPrev` and `cfNext` references are explicit in the *FlowGraph* metamodel, determining the corresponding next node `nextFlowNode` in the AST structure is more challenging and requires a non-trivial recursive search in the AST. In our approach, this search is specified using SDM or Java, and is integrated in the TGG rule via a virtual link referred to as a *binding expression*. In `AssignmentWithExpRule`, the binding expression is the link from `assignmentNode` to `nextFlowNode`, which is visualized as a dashed arrow leading to an object with a bold border. Using binding expressions, complex context relationships in TGG rules can be indicated by requiring a virtual link that does not really exist according to the corresponding metamodel. Our tool automatically creates a helper method, which is invoked when navigating along this link, and can be implemented with SDM or Java.

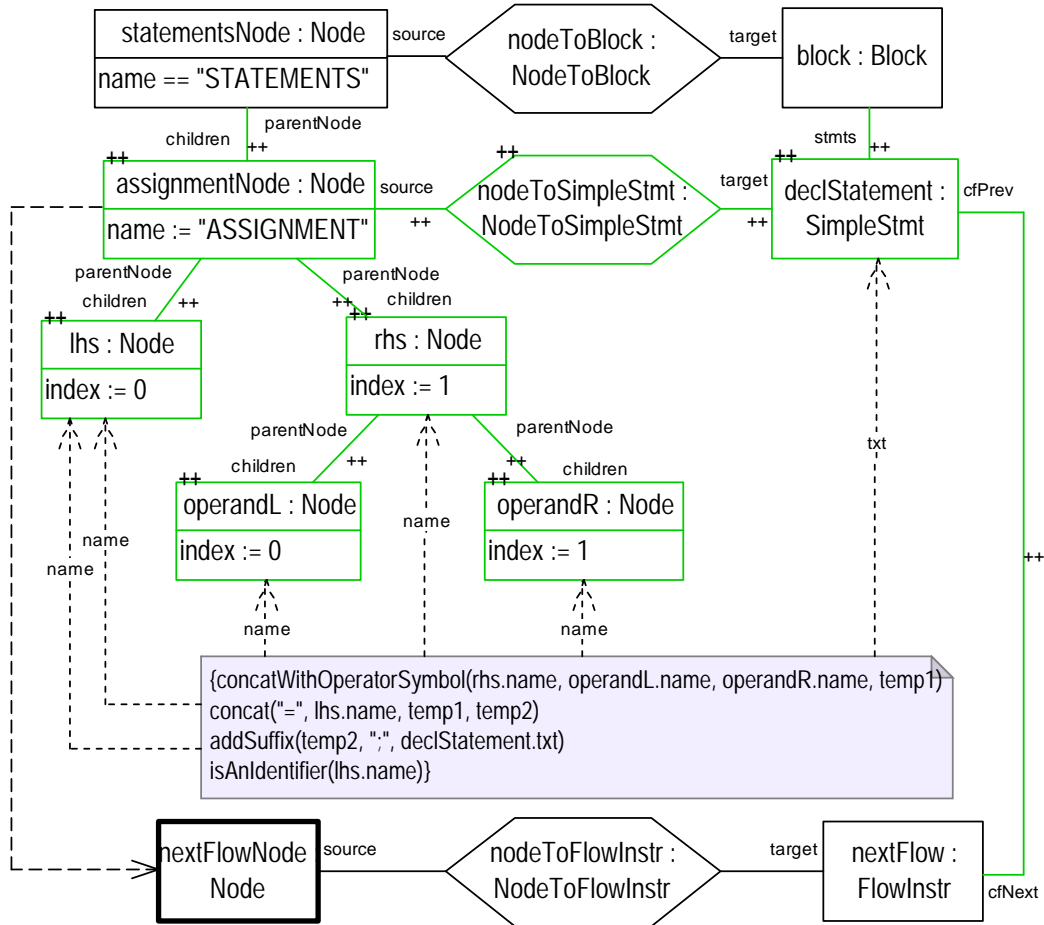


Figure 2: TGG rule AssignmentWithExpRule for creating assignment statements

Formally, binding expressions are equivalent to a separate pre-processing step in which all virtual links are created explicitly in an instance of an appropriately augmented metamodel. The TGG rules are then viewed as operating on this augmented metamodel. In practice, however, it is much more convenient to “find” these links *on-the-fly* as required. In this manner, binding expressions allow the controlled integration of SDM or Java code in TGG rules to determine complex context relationships.

As a final feature, every TGG rule can invoke a *post-processing method* (implemented with SDM or Java) to perform some final tasks that are difficult or impossible to specify directly in the TGG rule. In AssignmentWithExpRule, the index of the created assignmentNode has to be set correctly for the backward transformation, which requires a non-trivial recursive search in the tree structure due to if/else, loop, and return/break statements.

Post-processing methods are clearly a black-box extension and beyond any formal reasoning. In our practical experience, however, post-processing can be kept in most cases to a bare minimum by using a combination of attribute constraints and binding expressions.

3 Related Work

Related approaches can be divided into three main groups: (i) Other ways of increasing the expressiveness of TGGs, (ii) Other bidirectional approaches, and (iii) Achieving bidirectionality with unidirectional transformation languages.

(i) TGGs with a separated pre- and post-processing step: A justified question is if it is not better to use TGGs with a clearly separated pre-processing and post-processing step instead of our features for integrating auxiliary methods in the TGG rules. Our approach, however, avoids an extra traversal through the input and output models as *binding expressions* can be used as required to determine context relationships before a TGG rule is applied and a *post-processing method* to complete the rule after application. In this way, the traversal strategy applied by the TGG control algorithm can be used for *on-the-fly* processing as required. Our practical experience is that users tend to regard TGGs as not being worth the effort if complex pre-processing and post-processing with a separate traversal is necessary. With our approach, the TGG is more in focus and is used to structure the transformation in an iterative manner.

(ii) Other bidirectional languages and tools: For a detailed survey of bidirectional languages other than TGGs we refer to [14]. Compared to other bidirectional approaches, TGGs are advantageous as there exist multiple, actively maintained TGG implementations (interpretative, generative and hybrid approaches) with different strengths and weaknesses. Furthermore, most TGG implementations are Ecore/EMF/Eclipse compatible and are thus suitable in an MDE context together with other EMF tools and infrastructure. In a different context, however, other tools are probably more suitable, e.g., for bidirectional XML manipulation [10] and for bidirectional string and tree manipulation [3]. Depending on the application scenario, it might be more appropriate to derive a backward transformation from a given forward transformation [6] as opposed to describing the simultaneous build-up of model triples.

(iii) Combination of unidirectional transformation languages: A combination of unidirectional transformation languages as an alternative to a bidirectional language is advantageous for obvious reasons: Unidirectional transformation languages are typically better established, stable, expressive and have better tool support. We, however, regard transformation languages with explicit support for bidirectionality as superior as they enable a high-level specification. In the case of TGGs, the specified consistency relation can be used to automatically derive other useful *operational* transformations to support conformance testing, consistency checking and link creation for existing source and target models, and incremental model synchronization [9]. Supporting these scenarios and guaranteeing suitable formal properties [8] is quite challenging using separate forward and backward transformations.

4 Future Work

In the near future we will establish, compare, and experiment with various ideas for modularity concepts to improve the maintainability of TGG rules. Efficiency and scalability of transformations with TGGs are also crucial points that must be improved. We are currently using TGGs for an increasing number of non-trivial transformations including industrial projects together with Siemens AG and internally (i.e., as part of eMoflon) for a bootstrap of a textual syntax for Ecore/SDM/TGGs with TGGs. We aim to establish a transformation zoo for TGGs consisting of numerous representative examples.

References

- [1] Anthony Anjorin, Marius Lauder, Sven Patzina & Andy Schürr (2011): *eMoflon : Leveraging EMF and Professional CASE Tools*. In: *Proc. of Informatik 2011, LNI 192*, GI, p. 281.
- [2] Anthony Anjorin, Gergely Varró & Andy Schürr (2012): *Complex Attribute Manipulation in TGGs with Constraint-Based Programming Techniques*. In: *Proc. of Bx 2012, ECEASST 49*, EASST.
- [3] Aaron Bohannon, J.N. Foster, B.C. Pierce, Alexandre Pilkiewicz & Alan Schmitt (2008): *Boomerang: Resourceful Lenses for String Data*. *ACM SIGPLAN Notices* 43(1), pp. 407–419, doi:10.1145/1328897.1328487.
- [4] Thorsten Fischer, Jörg Niere, Lars Torunski & Albert Zündorf (2000): *Story Diagrams: A New Graph Rewrite Language Based on the Unified Modeling Language and Java*. In Hartmut Ehrig, Gregor Engels, Hans-Jörg Kreowski & Grzegorz Rozenberg, editors: *Proc. of TAGT 1998, LNCS 1764*, Springer, pp. 296–309, doi:10.1007/978-3-540-46464-8_21.
- [5] Florian Heidenreich, Jendrik Johannes, Sven Karol, Mirko Seifert & Christian Wende (2009): *Derivation and Refinement of Textual Syntax for Models*. In Richard F. Paige, Alan Hartman & Arend Rensink, editors: *Proc. of ECMDA-FA 2009, LNCS 5562*, Springer, pp. 114–129, doi:10.1007/978-3-642-02674-4_9.
- [6] Soichiro Hidaka, Zhenjiang Hu, Kazuhiro Inaba, Hiroyuki Kato & Keisuke Nakano (2011): *GRoundTram: An Integrated Framework for Developing Well-Behaved Bidirectional Model Transformations*. In Perry Alexander, Corina Pasareanu & John Hosking, editors: *Proceedings of ASE 2011, IEEE*, pp. 480–483, doi:10.1109/ASE.2011.6100104.
- [7] Tassilo Horn (2013): *The TTC 2013 Flowgraphs Case*. In: *Sixth Transformation Tool Contest (TTC 2013)*, *EPTCS* this volume.
- [8] Felix Klar, Marius Lauder, Alexander Königs & Andy Schürr (2010): *Extended Triple Graph Grammars with Efficient and Compatible Graph Translators*. In Andy Schürr, C. Lewerentz, G. Engels, W. Schäfer & B. Westfechtel, editors: *Nagl Festschrift, LNCS 5765*, Springer, pp. 141–174, doi:10.1007/978-3-642-17322-6_8.
- [9] Marius Lauder, Anthony Anjorin, Gergely Varró & Andy Schürr (2012): *Efficient Model Synchronization with Precedence Triple Graph Grammars*. In Hartmut Ehrig, Gregor Engels, Hans-Jörg Kreowski & Grzegorz Rozenberg, editors: *Proc. of ICGT 2012, LNCS 7562*, Springer, pp. 401–415, doi:10.1007/978-3-642-33654-6_27.
- [10] Christian Nentwich, Licia Capra, Wolfgang Emmerich & Anthony Finkelstein (2002): *Xlinkit: A Consistency Checking and Smart Link Generation Service*. *ACM Transactions on Internet Technology* 2(2), pp. 151–185, doi:10.1145/514183.514186.
- [11] Terence John Parr (2004): *Enforcing strict model-view separation in template engines*. *Proc. of WWW '04*, p. 224, doi:10.1145/988672.988703.
- [12] Terence John Parr (2007): *The Definitive ANTLR Reference: Building Domain-Specific Languages*. The Pragmatic Bookshelf.
- [13] Andy Schürr (1994): *Specification of Graph Translators with Triple Graph Grammars*. In G. Tinhofer, editor: *Proc. of WG 1994, Lecture Notes in Computer Science 903*, Springer, pp. 151–163, doi:10.1007/3-540-59071-4_45.
- [14] Perdita Stevens (2008): *A Landscape of Bidirectional Model Transformations*. In Ralf Lämmel, Joost Visser & João Saraiva, editors: *Proc. of GTTSE 2007, LNCS 5235*, Springer, pp. 408–424, doi:10.1007/978-3-540-88643-3_10.
- [15] Transformation Solution with eMoflon hosted on Share (2013): http://is.ieis.tue.nl/staff/pugorp/share/?page=ConfigureNewSession&vdi=XP-TUe_TGG-Comparison_eMoflon_TTC13_eMoflon.vdi.