



MASTER'S THESIS

Towards a Unifying Framework for Modelling and Executing Model Transformations

Ivo van Hurne

18th June, 2014

Supervisors: dr. L. Ferreira Pires
dr. C.M. Bockisch

UNIVERSITY OF TWENTE.

*Faculty of EEMCS
Department of Computer Science*

A Thesis Submitted in Partial Fulfilment
of the Requirements for the Degree of
MASTER OF SCIENCE IN COMPUTER SCIENCE
at the
UNIVERSITY OF TWENTE

Enschede, June 2014

Abstract

Model-driven engineering is a software engineering technique which relies heavily on the use of models. They are not just used as documentation, but actually define the system. Transformations on the models are described using model transformation languages. There are a lot of different model transformation languages available, all having a different approach to model transformation.

We show that these languages are actually not that different at all. Based on the analysis of a varied selection of transformation languages we define a small number of primitive transformation operations that can be used to describe model transformations written in any transformation language. We define our own primitive transformation language, using just these operations, and verify our analysis by implementing a few well-known transformation languages in our own language, including a language not considered in the initial analysis.

We design an interpreter for our primitive language and show that the execution of model transformations with our interpreter is on par with their original interpreter.

Acknowledgements

I would like to thank a few people for their support during the writing of this thesis. First of all, I would like to thank Luís Ferreira Pires and Christoph Bockisch for taking over the supervision of the project. Without their help it probably would not have been possible to complete this thesis. I would also like to thank my fellow students at the SE-lab for their valuable feedback and support.

I'm indebted to Edwin Vlieg and Joost Diepenmaat for giving me the final push necessary to finish the project. Finally, I'm grateful to my family and friends for their never-ending support.

Contents

List of Figures	xiii
List of Tables	xv
List of Algorithms	xvii
1. Introduction	1
1.1. Motivation	1
1.2. Objectives	2
1.3. Research Questions	2
1.4. Research Approach	2
1.5. Outline	3
2. Model Transformation Languages	5
2.1. Models	5
2.2. Eclipse Modeling Framework	5
2.3. Model Transformation	6
2.4. Model Transformation Pattern	7
2.5. Transformation Languages	7
2.5.1. Transformation Approach	7
2.5.2. Rule Application Strategy	8
2.5.3. Model Representation	9
2.5.4. Tracing	9
2.6. Languages to Consider	10
2.7. Conclusions	10
3. Transformation Execution Algorithms	11
3.1. Graph Transformation	11
3.2. Controlled Graph Transformation	12
3.2.1. GROOVE	12
3.2.2. Henshin	12
3.3. Term Rewriting	16
3.4. ATL	20
3.5. Analysis of Commonality and Variability	21
3.5.1. Rule Level	21
3.5.2. Transformation Level	21

3.6. Conclusions	23
4. Primitive Model Transformation Language	25
4.1. Language Definition	25
4.1.1. Model Navigation Operations	25
4.1.2. Collection Operations	26
4.1.3. Logical and Arithmetic Operations	26
4.1.4. Functions	26
4.1.5. Exception Handling	27
4.1.6. Transformation-level Operations	27
4.1.7. Rule-level Operations	27
4.1.8. Operations not included	28
4.2. Application of the Primitive Language	29
4.2.1. Graph Transformation	29
4.2.2. Henshin	30
4.2.3. Term Rewriting	31
4.2.4. Declarative ATL	31
4.2.5. Imperative ATL	31
4.3. Related Work	32
4.4. Conclusions	32
5. Transforming Model Transformations	35
5.1. Architecture	35
5.2. Interpreter Implementation	35
5.3. Transformation Language Implementation	37
5.4. Conclusions	40
6. Model Transformation Scenarios	41
6.1. Common Use Cases	41
6.2. Object-oriented Class to Relational Table	42
6.2.1. Implementation	42
6.2.2. Results	42
6.3. Pull Up Class Attribute to Superclass	44
6.3.1. Implementation	44
6.3.2. Results	44
6.4. Conclusions	44
7. Conclusions and Future Work	47
7.1. A Common Execution Environment	47
7.2. Execution in the Common Environment	48
7.3. Limitations of the Common Environment	48
7.4. Looking ahead	49

A. Primitive Representations of the Algorithms	51
A.1. Single-Pushout Graph Transformation	51
A.2. Henshin	53
A.3. Stratego	56
A.4. Declarative ATL	57
A.5. Imperative ATL	59
B. Thrascias Implementations of the Algorithms	61
B.1. SimpleGT	61
B.2. Declarative ATL	64
C. Implementations of the Mapping Scenario	67
C.1. SimpleGT	67
C.2. Declarative ATL	70
D. Implementations of the Refactoring Scenario	73
D.1. SimpleGT	73
D.2. Declarative ATL	75

List of Figures

2.1. Kernel of Ecore model [1]	6
2.2. Model Transformation Pattern	7
2.3. Model transformation pattern (including metamodel) [2]	8
3.1. Simple graph transformation example	13
3.2. Henshin transformation units [3]	18
5.1. Model Transformation Pattern	36
5.2. Transformation engine architecture	36
5.3. Thrascias Trace structure	37
5.4. Thrascias abstract syntax	38
6.1. Metamodels for mapping scenario [4]	42
6.2. EMF 'extlibrary' metamodel (adapted from [5])	43
6.3. 'extlibrary' model after Pull Up Class Attribute transformation	43
6.4. Model for refactoring scenario	44

List of Tables

3.1. Language constructs in GROOVE control language	15
3.2. A selection of Stratego language constructs	19
3.3. Comparison of transformation operations at the rule level.	23
3.4. Comparison of transformation operations at the transformation level. .	24
4.1. Primitive transformation language operations	29
4.2. Comparison of transformation primitives	33
5.1. Thrascias operations	39

List of Algorithms

3.1. Pseudo-code representation of single-pushout graph transformation. . .	14
3.2. Pseudo-code representation of Henshin's execution algorithm.	17
3.3. Pseudo-code representation of the Stratego rule application algorithm. .	19
3.4. Pseudo-code representation of Stratego's cascading transformation strategy.	20
3.5. Pseudo-code representation of the ATL transformation algorithm. . . .	22

1

Introduction

1.1. Motivation

The *Model-Driven Architecture* (MDA) [6] was introduced in 2001 by the Object Management Group (OMG) to aid the integration, interoperability and evolution of software systems. MDA is an approach to specifying software systems that separates specification of functionality from its implementation on a specific technology platform.

Kent [7] introduced the term *Model-Driven Engineering* (MDE), which in a sense extends MDA. MDE is not only concerned with the separation of platform-independent from platform-specific, but also with the development process and extra dimensions like versioning. Contrary to MDA, MDE is not necessarily based on OMG standards. Instead it uses general standard-independent modelling concepts.

MDE relies heavily on the use of models. They are not just used as documentation but define the system at a high abstraction level, omitting any technology-specific information. A model can describe, for example, the structure or the behaviour of a system. The actual implementation of the system can be generated from such a model via a model transformation. Transformations can also be used for reverse engineering (implementation to model) and the transfer of information between systems.

In MDE, model transformations are usually defined using a *model transformation language*. Research has led to the creation of several transformation languages and at present new languages continue to appear.

New model transformation languages can be useful to test and demonstrate novel ideas, and to specify languages tailored for a given domain. However, the development of these languages is a time-consuming process, as most of them are implemented from scratch and as such they are difficult to extend and adapt.

1.2. Objectives

Although the languages all have their own particular approach and features, they have the same main purpose. This means they are likely to share some common features. For example, in all transformation languages, models can be manipulated and navigated in some way (e.g. using OCL [8]). Languages often share common structures like transformation rules, and may provide control over the execution order of these rules.

The objective of this work is to create a unifying framework for executing transformation languages that exploits the commonalities between a set of languages, while being able to execute each one of them. This would enable faster prototyping of new ideas, easier reuse of (parts of) transformations and the composition of transformations written in different languages.

1.3. Research Questions

How can we execute different model transformation languages in a common environment, while retaining the variability between the languages?

This question has been decomposed into three subquestions:

1. What should a common execution environment for model transformation languages look like?
2. How do we build support for executing different transformation languages in this common environment?
3. What are the limitations of executing transformation languages in a common environment?

1.4. Research Approach

- We first take a look at the principles of models and model transformation languages. Out of the vast amount of transformation languages available, we select a few representative languages for a more detailed analysis.
- Second, we investigate the execution algorithms of these languages. We look into the commonalities and variabilities, which results in a set of primitive transformation operations.
- Third, we define a language for specifying transformation languages in our common environment.
- Fourth, we implement an interpreter for our language in order to execute the specified languages.
- Finally, we present a number of test cases to validate our framework and identify its limitations.

1.5. Outline

This thesis is organized as follows. Chapter 2 introduces model transformation and transformation languages. Chapter 3 analyses the commonalities and variabilities between the languages, and take a look at their execution algorithms. Chapter 4 defines our language for specifying transformation languages in our framework. Chapter 5 introduces our implementation of the framework and its architecture. Chapter 6 discusses test cases for validating the implementation and its limitations. Chapter 7 draws conclusions about the results and sketches future work.

2

Model Transformation Languages

We start with an introduction to model transformation and provide an overview of existing transformation languages.

2.1. Models

We mentioned before that models are the cornerstone of MDE. The word *model* is used to describe lots of different concepts in different contexts. We will use the definition by Kurtev [9], which summarizes some of the existing definitions: “A model represents a part of the reality called the object system and is expressed in a modelling language. A model provides knowledge for a certain purpose that can be interpreted in terms of the object system.”

Most models conform to a *metamodel*, which is a model of a modelling language and basically describes its abstract syntax. Because a metamodel is itself a model, it is an instance of another metamodel (also known as a *meta-metamodel*).

2.2. Eclipse Modeling Framework

The Eclipse Modeling Framework (EMF) is a framework for modelling and code generation of applications based on structured data models [5]. In EMF, metamodels are defined as instances of the *Ecore model*. Additionally, the Ecore model itself is defined as an Ecore model.

Figure 2.1 shows the main part of the Ecore model [1]. EClasses are used to define classes along with their attributes (EAttribute) and references (EReference). EPackages can be used to group the classes.

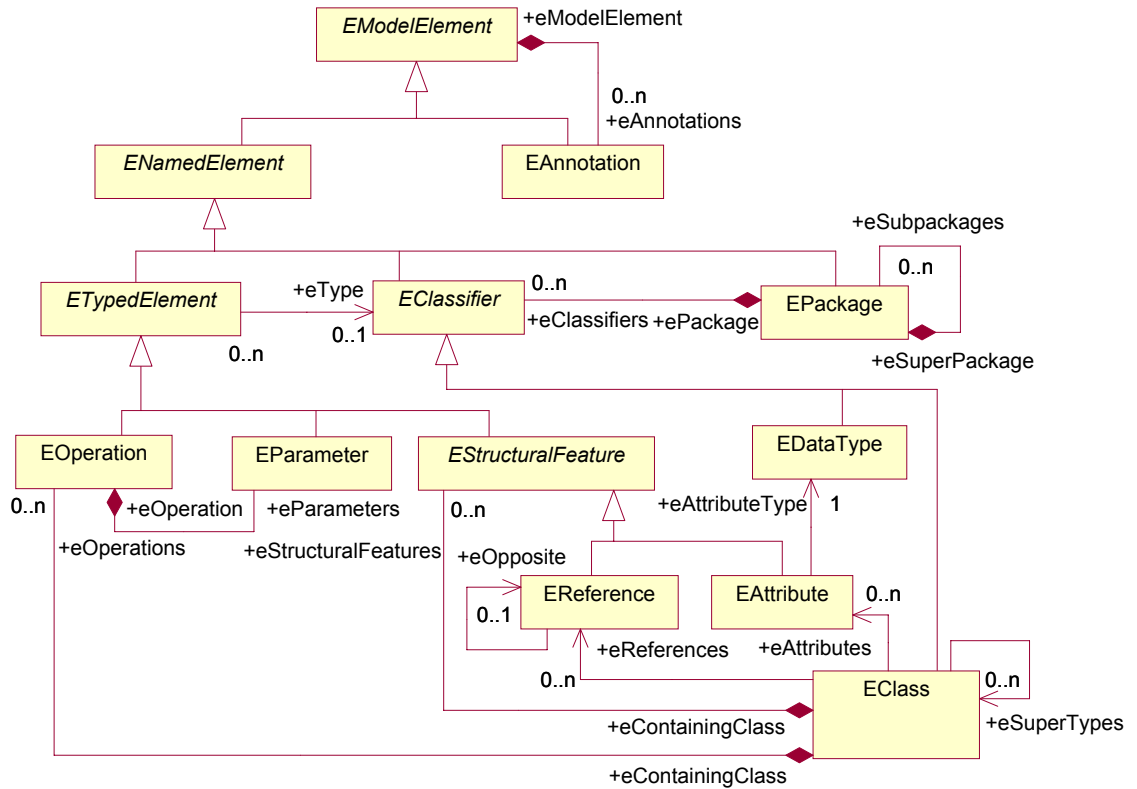


Figure 2.1.: Kernel of Ecore model [1]

EMF provides basic Java code generation for its models, including a simple editor, support for basic model modification operations and persistence.

2.3. Model Transformation

In the Object Management Group’s MDA Guide, model transformation is defined as “the process of converting one model to another model of the same system” [6]. However, this definition does not include the specification of a transformation and unnecessarily restricts transformations to models of the same system. Additionally, it does not take into account transformation to and from multiple models.

We therefore use the definition by Mens and van Gorp [10]: “A model transformation is a process of automatic generation of one or multiple target models from one or multiple source models, according to a transformation definition, which is expressed in a model transformation language.”

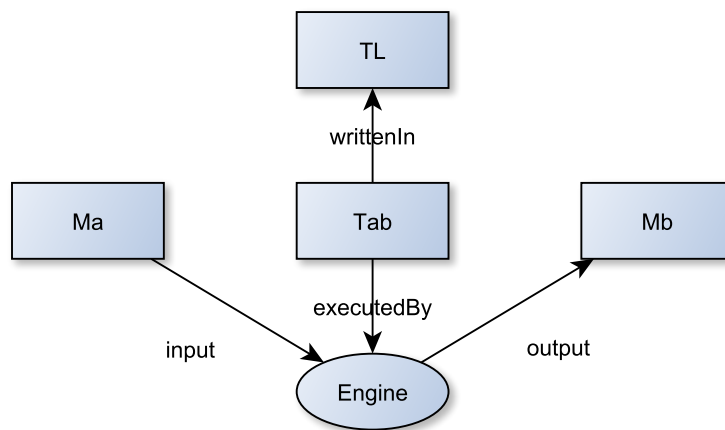


Figure 2.2.: Model Transformation Pattern

2.4. Model Transformation Pattern

Model transformations can be described using the pattern shown in Figure 2.2. We have a source model M_a , a target model M_b and a transformation definition T_{ab} . The transformation definition is written in a transformation language TL . When executed, T_{ab} transforms M_a to M_b .

As mentioned before multiple source and target models can be used in the same transformation, extending the pattern accordingly.

As a model is an instance of a metamodel, we can extend the pattern to include the metamodels (Figure 2.3). Some transformation languages do not need metamodels because they operate on generic graphs or ASTs (e.g. GROOVE, Stratego). Other languages do need metamodels.

2.5. Transformation Languages

There are many different model transformation languages available, but they can nevertheless be classified into a number of categories. Because we are mainly interested in the execution algorithms of model transformation languages, we consider four aspects for classification: *transformation approach*, *rule application strategy*, *model representation* and *tracing* [11].

2.5.1. Transformation Approach

Mens and van Gorp [10] state that the most important distinction between the languages is whether they use a *declarative* or an *operational* approach.

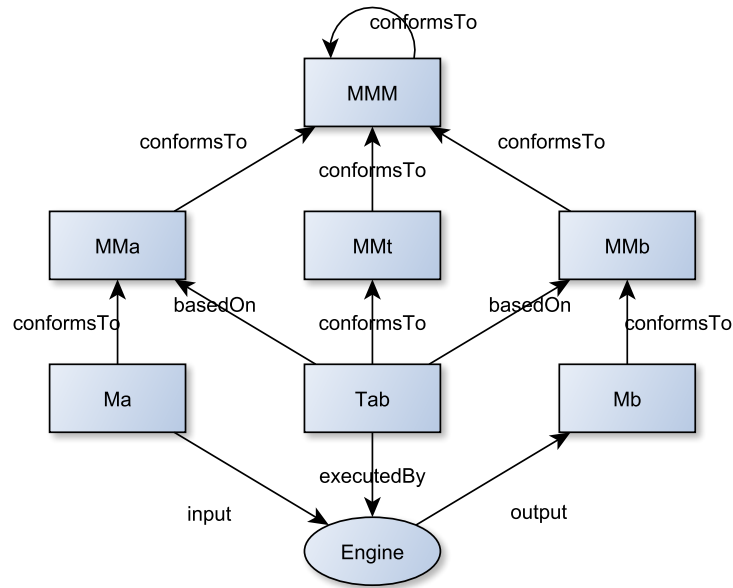


Figure 2.3.: Model transformation pattern (including metamodel) [2]

Declarative Approach The declarative approach is also known as the *relational* approach. Developers specify relations between elements of the source and target models. The transformation language has to figure out the necessary transformation steps by itself. Examples of declarative languages are Henshin [3] and QVT Relations [12].

Imperative Approach The imperative approach is also known as the *operational* approach. Developers specify the exact operations that have to be performed to transform source models to target models. An example of an imperative transformation language is QVT Operational Mappings [12].

Hybrid Approach A third approach has yet to be mentioned: the *hybrid* approach. Strictly speaking this is not an approach separate from declarative or imperative, but simply makes both available to the developer in a single language. An example of a hybrid language is ATL [13].

2.5.2. Rule Application Strategy

A transformation usually consists of a number of units. In most languages these units are rules [11]. Rules can appear in many guises, for example a rewrite rule with a left-hand side and a right-hand side; or a function that takes an input pattern and produces output by using an expression.

The rules in a transformation can often be applied in different orders. Furthermore, rules can match at multiple model regions at the same time. Transformation languages

therefore use a strategy to choose the order in which transformation rules should be applied and at which model region. There are two kinds of strategies [11]:

Deterministic A deterministic strategy can either be *explicit* or *implicit*. In the first case the application order is defined by an explicit strategy, possibly separate from the transformation rules themselves (example: Stratego [14]).

In the second case the developer has but indirect control over the application order. The result of the transformation is deterministic, but can be changed by changing the matching patterns and logic of the rules (example: ATL [13]).

Non-deterministic In a non-deterministic strategy a rule and/or application location is chosen randomly from matching rules or model regions (example: graph transformation languages [15]). It is important to note that even if the application strategy is non-deterministic, the result of the transformation can still be deterministic, in case all the different application orders result in the same output model.

2.5.3. Model Representation

Models used by transformation languages can be represented in various ways. Some possibilities are graphs (e.g. graph transformation), abstract syntax trees (e.g. Stratego) and Eclipse Modelling Framework models [16] (e.g. Henshin, ATL).

2.5.4. Tracing

Traces record information about the execution of a transformation. Commonly traces are used to map source elements to target elements [11]. They can be used for analysing how changes in the source model affect the target models, for synchronizing models and for debugging.

There are various types of traces:

- Different kinds of information can be recorded, such as, the rule that created a trace or the time of creation.
- Information can be recorded at different abstraction levels, for example, only for top-level transformations.
- Information can be recorded for different scopes, for example, only for particular transformation rules.
- Traces can be stored in a number of locations, such as, in the source model, target model or in a separate location.

Some transformation languages provide support for tracing in the language itself, or even create traceability links automatically (ATL, QVT). In other languages tracing has to be done manually as part of the transformation definition.

2.6. Languages to Consider

We choose to analyse five declarative model transformation languages. Based on the classification above, they have significantly different characteristics and they are well-known representatives of their kind.

Graph transformation Not really a language as such, but rather a class of transformation languages with graph-based models, a non-deterministic rule application strategy and no inherent tracing support.

GROOVE Graph transformation language extended with explicit deterministic rule application strategy. No inherent tracing support.

Henshin Eclipse modelling framework-based language with explicit deterministic rule application strategy. No inherent tracing support.

Stratego Abstract syntax tree-based term-rewriting language with explicit deterministic rule application strategy. No inherent tracing support.

ATL EMF-based language with implicit deterministic rule application strategy and inherent tracing support. Although ATL is a hybrid language we only consider its declarative part.

We discuss these languages in detail in Chapter 3. We choose only declarative languages, so we can validate our analysis later on using a language with a different transformation approach (imperative ATL).

2.7. Conclusions

We have briefly introduced models and model transformation concepts. We have shown that there are many model transformation languages with different characteristics. Furthermore we have explained our choice of languages for analysis. Next, we discuss the execution algorithms of the chosen languages.

3

Transformation Execution Algorithms

In Chapter 2 we have selected five transformation languages for further analysis: graph transformation, GROOVE, Henshin, Stratego and declarative ATL. In this chapter we take a look at the way they are executed and try to find commonalities between them.

3.1. Graph Transformation

The first execution algorithm we consider is the algorithm used in graph transformation. In this kind of transformation a model is described as a graph. There are multiple ways to do this, the most simple being the mapping of entities to nodes and relations to edges. However, once more complicated relations are used (e.g. multiplicity) more complicated mappings are necessary [17]. This, however, is outside the scope of this thesis.

To transform a model, graph rules can be defined that specify pre-conditions for the application of a rule (called *left-hand side*) and post-conditions that have to be satisfied after the application of a rule (*right-hand side*) [15].

There are actually a number of different approaches to graph transformation that can be used. The way in which new elements are added and old elements are removed is slightly different for these approaches, but the general idea is the same. In this thesis we illustrate graph transformation using the *single-pushout approach* (SPO).

Algorithm 3.1 describes the approach in pseudo-code. We first select an arbitrary transformation rule. We try to find an arbitrary occurrence of its left-hand side L (Figure 3.1a) in a graph G (Figure 3.1b). This happens in line 10 of Algorithm 3.1 using `G.getMatches()`. During the matching that is taking place here, any possible negative application conditions are also taken into account. Negative application conditions are conditions under which a rule should not be applied. For example, one could add

a condition to the rule in Figure 3.1a that prevents guests from occupying multiple rooms.

Once matches have been found, we iterate over them in a non-deterministic order. All graph elements that exist in L but not in R are then deleted from G (lines 12-14). Subsequently, we check for any dangling edges in G , and delete them (lines 16-19). Finally in lines 21-23, we add the objects that exist in R but not in L to G (Figure 3.1c) [18].

If a transformation rule has been applied successfully, the transformation is restarted and we again start to look for matches for an arbitrary transformation rule. If no matches were found for a transformation rule we try a different rule. The transformation ends when none of the transformation rules can be applied.

Because of the non-deterministic choice of rules there is no guaranteed confluence. This means executing the transformation multiple times on the same model does not necessarily yield the same result. Moreover, termination of a transformation is not guaranteed and a rule might not be executed even if there exists a match for this rule in the graph.

3.2. Controlled Graph Transformation

In a graph transformation, rules are self-contained units that specify the exact pre-conditions for their application [19]. Rule ordering is often a major reason for specifying these conditions. Because the conditions can be very complex and rules can implicitly depend on other rules, the transformation is not always easy to understand. A possible solution is to move the complex pre-conditions out of the rules, specifying them using so-called control expressions. In other words: an explicit deterministic rule application strategy is introduced.

3.2.1. GROOVE

Staijen has defined a language for specifying such strategies for the GROOVE graph-transformation toolkit [20].

Table 3.1 lists the language constructs available in this language. The constructs take one or more expressions, which can be either rules or other constructs.

3.2.2. Henshin

Another language for graph transformations with explicit rule application strategy is Henshin [21], developed by Arendt et al. [3]. Henshin operates on EMF models, but the models are represented as graphs.

The application strategy is defined using *transformation units* (Figure 3.2). A rule is itself the most basic transformation unit. Other kinds of transformation units can have subunits that are executed in a particular order, or under particular circumstances. A transformation unit can be applicable, which means it will be executed during the transformation. Arendt et al. discuss some of these transformation units in their articles:

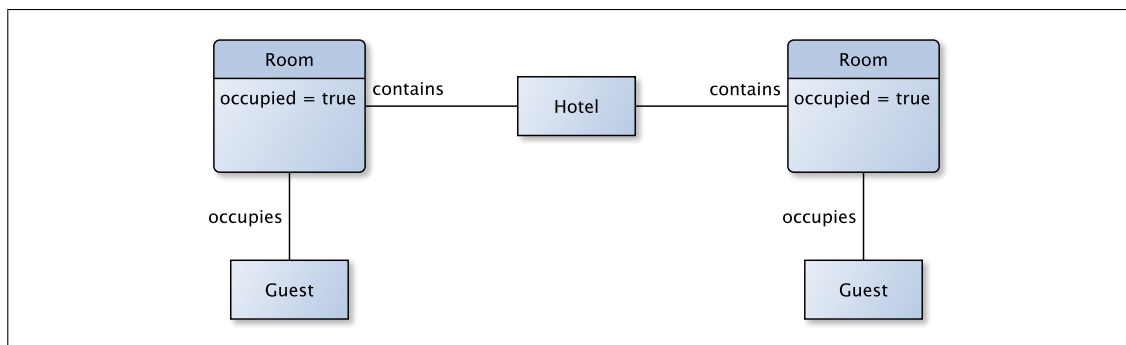
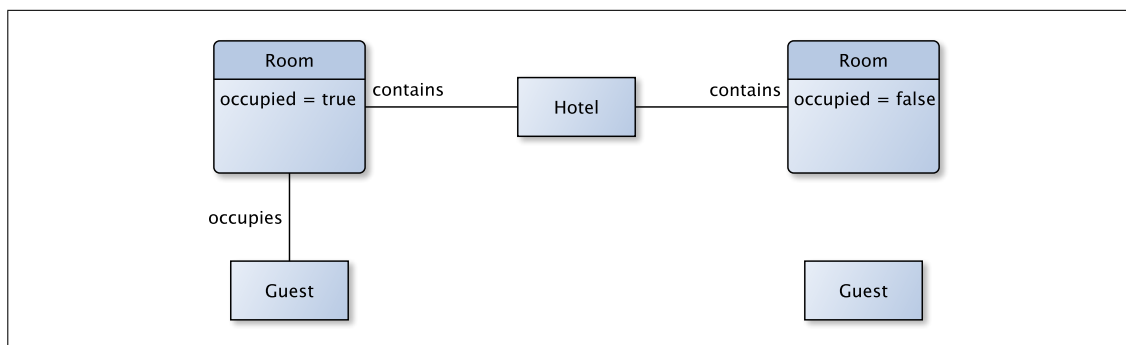
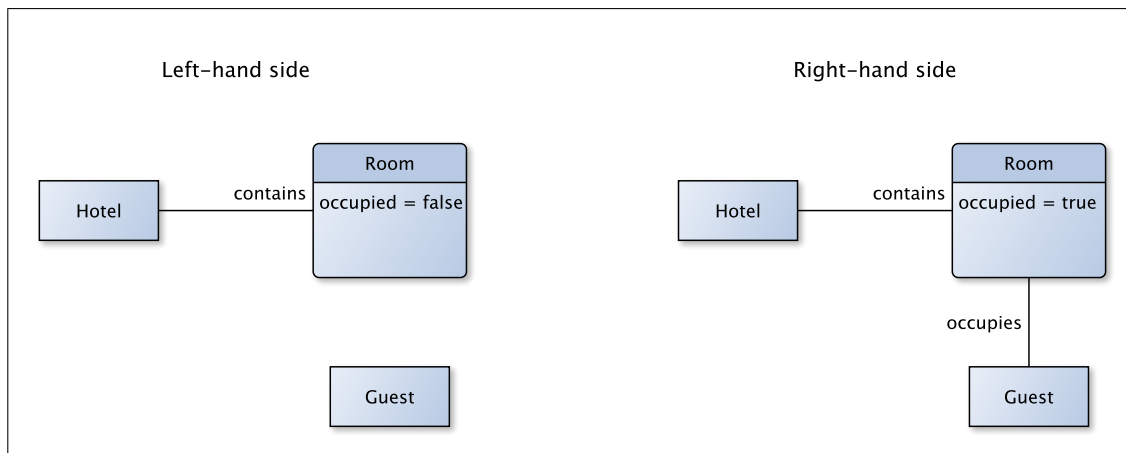


Figure 3.1.: Simple graph transformation example

Algorithm 3.1 Pseudo-code representation of single-pushout graph transformation.

```

1 InputModel G
2 Transformation transformation
3
4 // Continue until no application is possible
5 while True:
6     boolean ruleApplied = false
7
8     forall rule in transformation.getRules():
9         // Find occurrences of rule in graph
10        forall oL in G.getMatches(rule.getL()):
11            // Find elements in  $L \setminus R$ 
12            Set deletedElements = oL - rule.getR()
13            // Delete elements from G, creating D
14            Model D = G - deletedElements
15            // Check for dangling edges
16            forall e in D.edges():
17                if e.source == null || e.target == null:
18                    // Delete edge if dangling
19                    D = D - e
20            // Find elements in  $R \setminus L$ 
21            Set addedElements = rule.getR() - oL
22            // Glue new elements to D, creating new G (H)
23            G = D + addedElements
24            ruleApplied = true
25            break
26
27        if ruleApplied:
28            // rule applied, start from beginning
29            break
30
31    if !ruleApplied:
32        // no rules left, stop execution
33        break

```

Table 3.1.: Language constructs in GROOVE control language

Construct	Description
<i>ruleName</i>	Execute a rule.
true	Behaves like a rule that is always successful and does not change the underlying structure.
E1 E2	Non-deterministic choice. Execute either E1 or E2.
E1 ; E2	Sequential composition. First execute E1, then E2.
E*	Execute E an arbitrary number of times.
alap E	Repeat the execution of E as long as it applies.
try E1	Execute E1, skip if it does not apply.
try E1 else E2	First try to execute E1, then execute E2 only if E1 fails.
if (E1) E2	First execute E1, then afterwards E2 if E1 succeeds.
if (E1) E2 else E3	First execute E1, then afterwards E2 if E1 succeeds. If E1 fails, execute E3.
while (E1) do E2	Execute E1 and afterwards E2, and again E1 until the execution of E1 fails.
until (E1) do E2	Try to execute E1, then execute E2 and again E1 if E1 fails.

Rule A transformation rule [3]. It is applicable if there is a match for its left-hand side in the model.

IndependentUnit Non-deterministic execution of its subunits. An `IndependentUnit` is always applicable. Subunits may be applied repeatedly. The `IndependentUnit` terminates if no applicable subunit remains [3].

PriorityUnit Execution of a subunit that is applicable and has the highest priority. A `PriorityUnit` is always applicable. Subunits may be applied repeatedly. The `PriorityUnit` terminates if no applicable subunit remains [22].

SequentialUnit Execution of its subunits in a predefined sequence. A `SequentialUnit` is only applicable if all subunits are applicable in the given order. The `SequentialUnit` terminates if all subunits terminate [3].

CountedUnit Execution of its subunit a specified number of times. It is only applicable if its subunit is applicable the given number of times. It terminates if the subunit terminates [22].

ConditionalUnit Conditional execution using references to subunits called `if`, `then`, and `else`. It terminates if its subunits terminate [3].

AmalgamationUnit Allows for the definition of multiple rules with common parts. It consists of a *kernel rule* and a set of *multi-rules* [3].

- The kernel rule is matched only once and serves as a common partial match for each multi-rule.
- The multi-rules are matched as often as possible.

Algorithm 3.2 describes the execution algorithm for Henshin in pseudo-code. Henshin supports the passing of parameters between transformation units, but we omit this in the pseudo-code.

3.3. Term Rewriting

The next transformation language we discuss is called Stratego. It was developed by Visser [14] and is based on term rewriting. Instead of graphs, term-rewriting transformations use *abstract syntax trees* (AST) to represent the model. Abstract syntax trees consist of *terms*, i.e. applications $C(t_1, \dots, t_n)$ of a constructor C to terms t_i , lists, strings or integers.

An AST can be transformed by replacing its terms with other terms in a manner specified by term rewriting rules. However, as with graph transformation, a term rewriting system does not necessarily terminate, or yield the same result if rules are applied in a different order (*confluence*). Therefore, extra rules are often added to specify where rules should be applied and in what order.

Algorithm 3.2 Pseudo-code representation of Henshin's execution algorithm.

```

1 Transformation transformationSystem
2
3 forall unit in transformationSystem:
4     processUnit(unit)
5
6
7 function processUnit(unit):
8     if unit instanceof Rule:
9         Match match = unit.findMatch()
10        if match != null:
11            unit.applyRule(match)
12
13    else if unit instanceof IndependentUnit:
14        // Works in the same way as the
15        // basic graph transformation algorithm
16
17    else if unit instanceof PriorityUnit:
18        TransformationUnit nextUnit = unit.getNextSubUnit()
19        boolean success = false
20        while nextUnit != null && !success:
21            success = processUnit(nextUnit)
22            nextUnit = unit.getNextSubUnit()
23
24    else if unit instanceof SequentialUnit:
25        TransformationUnit nextUnit = unit.getNextSubUnit()
26        boolean success = true
27        while nextUnit != null && success:
28            success = processUnit(nextUnit)
29            nextUnit = unit.getNextSubUnit()
30
31    else if unit instanceof CountedUnit:
32        int counter = unit.getCount()
33        boolean success = true
34        while counter > 0 && success:
35            success = processUnit(unit.getSubUnit())
36            counter--
37
38    else if unit instanceof AmalgamationUnit:
39        Match kernelMatch = unit.getKernelRule().findMatch()
40        if kernelMatch != null:
41            if unit.getKernelRule().applyRule(kernelMatch):
42                forall multiRule in unit.getMultiRules():
43                    forall match in multiRule.findMatch():
44                        multiRule.applyRule(match)
45
46    else if unit instanceof ConditionalUnit:
47        boolean success = processUnit(unit.getIf())
48        if success:
49            processUnit(unit.getThen())
50        else if unit.getElse() != null:
51            processUnit(unit.getElse())

```

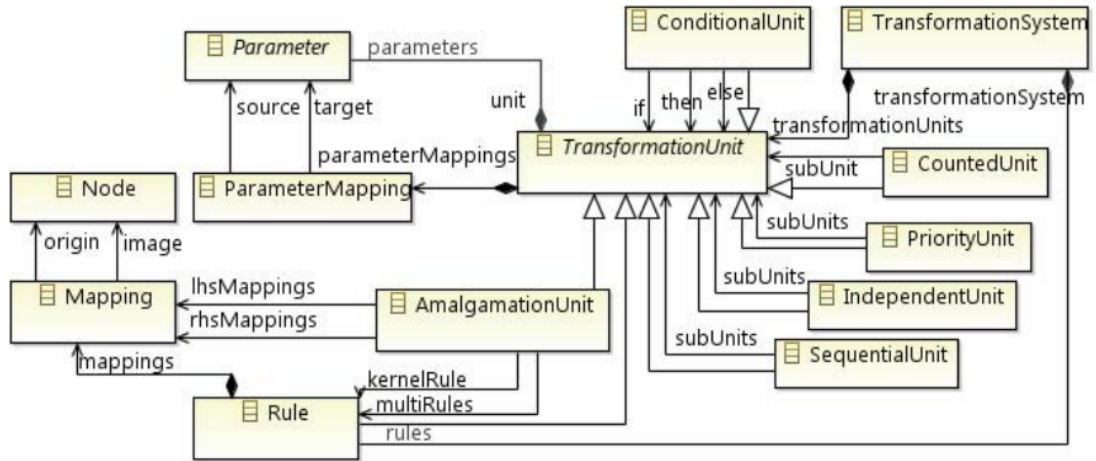


Figure 3.2.: Henshin transformation units [3]

In Stratego the rule application strategy is specified explicitly using higher-order rewriting rules. That is, the developer can specify rules that combine other rules and execute them in a particular order.

Stratego has a number of language constructs to facilitate this, described in Table 3.2.

Algorithm 3.3 describes Stratego’s rule execution algorithm in pseudo-code. First a rule is matched against the source model and its variables are bound to specific terms. Afterwards any possible rule preconditions are checked and additional variables are bound. Finally the rule is applied, transforming the model.

Visser provides a number of transformation idioms to illustrate the rule application strategies that could be expressed using Stratego. We take a look at one of them, namely cascading transformation.

Cascading transformation is the most basic strategy for term rewriting (Algorithm 3.4). Small independent transformations are cumulatively applied in order to reach a desired result. The strategy tries to apply any of the rules, starting at the bottom of the abstract syntax tree. Each successful application restarts this process. If no rule can be applied, the algorithm moves to a higher level in the tree. This continues until we reach the root and there are no rules left that apply.

Listing 3.1 shows how this strategy can be expressed in Stratego notation. We first traverse the tree to get to the bottom level using `bottomup(s)`. We try to apply the sequence of rules R_1 to R_n until one of the rules succeeds. If a successful rule application has taken place, `innermost(s)` restarts the transformation on the current subtree. Once no applicable rules remain, `bottomup(s)` moves us to the next level of the tree.

Table 3.2.: A selection of Stratego language constructs

Construct	Description
<code>ruleName : l -> r</code> <code>where s</code>	Define a rewrite rule with label <code>ruleName</code> , left-hand side <code>l</code> and right-hand side <code>r</code> . Optionally, preconditions <code>s</code> for the rule can be specified by adding a <code>where</code> part to the rule.
<code>ruleName</code>	Execute the rewrite rule labelled <code>ruleName</code> .
<code>E1 ; E2</code>	Sequential composition. Execute <code>E1</code> and after that <code>E2</code> .
<code>E1 <+ E2</code>	Deterministic choice. First try <code>E1</code> , execute <code>E2</code> only if <code>E1</code> fails.
<code>E1 + E2</code>	Non-deterministic choice. The same as <code><+</code> , but the first expression to try is chosen randomly.
<code>E1 < E2 + E3</code>	Guarded choice. If <code>E1</code> succeeds, execute <code>E2</code> . If <code>E1</code> fails, execute <code>E3</code> .
<code>where(E)</code>	Test whether <code>E</code> applies, but ignore the result of the application.
<code>not(E)</code>	Negation. Succeeds if <code>E</code> fails to apply.
<code>id</code>	Identity. Always succeeds with the original term as the result.
<code>all(E)</code>	Apply <code>E</code> to each direct subterm.
<code>one(E)</code>	Apply <code>E</code> to one direct subterm.
<code>try(E) = E <+ id</code>	Try to execute <code>E</code> , continue if the execution fails.
<code>repeat(E) = try(E;</code> <code>repeat(E))</code>	Repeat the execution of <code>E</code> until it fails.

Algorithm 3.3 Pseudo-code representation of the Stratego rule application algorithm.

```

1 function applyRule(Rule rule, AST tree):
2   // check for match and collect variable bindings
3   Set bindings = rule.l.match(tree)
4   // check extra preconditions
5   if rule.s != null:
6     Set precondition_bindings = rule.s.match(tree)
7   if rule.s == null || precondition_bindings != null:
8     bindings += precondition_bindings
9     if bindings != null:
10      // replace variables in term and replace tree
11      tree = rule.r.replaceVars(bindings)

```

Algorithm 3.4 Pseudo-code representation of Stratego’s cascading transformation strategy.

```

1 // cascading transformations
2 function cascade(Set rules, AST tree):
3     // traverse bottom-up
4     forall child in tree.getChildren():
5         cascade(rules, child)
6
7     // try to apply rules
8     forall rule in rules:
9         if applyRule(rule, tree):
10            // successful application => start applying all
11            // rules again bottom-up
12            cascade(rules, tree)
13            break

```

Listing 3.1: Cascading transformation expressed in the Stratego language. [14]

```

1 bottomup(s) = all(bottomup(s)); s
2
3 innermost(s) = bottomup(try(s; innermost(s)))
4
5 simplify = innermost(R1 <+ ... <+ Rn)

```

3.4. ATL

The last transformation language we discuss in this chapter is ATL [13]. Although ATL is a hybrid language we only consider its declarative part here. As explained in Chapter 2, we use the imperative part for validation purposes in Chapter 4.

Declarative ATL has three kinds of transformation rules: *matched*, *lazy* and *unique lazy*.

Matched rules Matched rules match a specific type of source elements, possibly further restricted by a guard condition. A developer specifies the target elements that have to be created from these source elements, and how their properties should be initialized. Each source element may only be matched by a single matched rule.

ATL executes matched rules automatically in an arbitrary order. It also creates traceability links between source and target elements.

Lazy rules Lazy rules are matched rules that are not executed automatically. They can be triggered by other rules as many times as necessary. Traceability links are not created for these rules.

Unique lazy rules Unique lazy rules are lazy rules that are applied at most once for a specific match. If triggered more than once, they return the target elements that were already created earlier.

Algorithm 3.5 describes ATL's transformation execution algorithm in pseudo-code. First, rules are matched and a list of matches is created. Second, the target elements are created. Last, the properties of the target elements are initialized.

Although we do not explicitly include the execution algorithm for (unique) lazy rules here, the algorithm for matched rules sans the matching can easily be reused for those kinds of rules.

3.5. Analysis of Commonality and Variability

Now that we have introduced the transformation execution algorithms of the various transformation languages, we proceed by analysing their commonalities and variabilities. Looking at the execution of a model transformation, we can distinguish two levels at which the execution takes place: the rule level and the transformation level.

3.5.1. Rule Level

At the rule level we look at the operations that take place during rule execution. As we are considering declarative languages, developers do not explicitly specify most of these operations, i.e. they happen 'behind the scenes'.

The transformation operations we identified are listed in Table 3.3. In every language the source model is navigated and elements from the source patterns are matched. Creating elements and setting properties is also something that every language is capable of.

In contrast, deleting elements or properties is not always possible. In ATL it is not possible to navigate the target model, and consequently it is not possible to delete parts of this model. Deletion of elements is possible when using ATL's *refining mode*, but this mode can only be used if source and target models conform to the same metamodel. Moreover, some language features are disabled in this mode.

ATL does have two operations that are not available in other languages, namely for creating and querying a trace. Traces can be created in other languages, but are not supported natively in the implementation of these languages. Instead, they have to be specified explicitly in the transformation definition using a custom-made metamodel.

3.5.2. Transformation Level

At the transformation level we look at operations that make up a rule application strategy. The operations we identified are listed in Table 3.4.

The most basic operation in such a strategy is executing a rule. In some languages (e.g. Stratego) it is possible to check if there are matches for a rule in the source model,

Algorithm 3.5 Pseudo-code representation of the ATL transformation algorithm.

```

1 Transformation transformation
2 InputModel inputModel
3 List matches
4 List targetElements
5 Dictionary traces
6
7 // match rules
8 forall rule in transformation.getMatchedRules():
9     forall match in inputModel.getMatches(rule):
10         matches += match
11
12 // execute matched rules
13 forall match in matches:
14     // create target elements
15     forall t in match.getRule().getTargets():
16         Object target = new (t.getType())()
17         targetElements += target
18         // create trace links
19         traces[match.getSource()] = target
20
21 // initialize target-element properties
22 forall target in targetElements:
23     forall property in target.getProperties():
24         // if primitive, assign directly
25         if isPrimitive(property.getContent()):
26             target.setProperty(property.getName(), property.getContent())
27         else:
28             // if source element, resolve, then assign
29             Trace trace = traces[property.getContent()]
30             if trace != null:
31                 target.setProperty(property.getName(), trace)
32
33             // if trace query, resolve, then assign
34             else if isQuery(property.getContent()):
35                 Trace trace = traces[property.getContent().getSource()]
36                 target.setProperty(property.getName(), trace)
37
38             // if target element, assign directly
39             else:
40                 target.setProperty(property.getName(), property.getContent())

```

Table 3.3.: Comparison of transformation operations at the rule level.

Operations	GT	GROOVE	Henshin	Stratego	ATL
Navigate model	X	X	X	X	X
Match pattern	X	X	X	X	X
Create element	X	X	X	X	X
Delete element	X	X	X	X	
Set property	X	X	X	X	X
Delete property	X	X	X	X	
Create trace					X
Query trace					X

but not actually execute the rule. ATL is a special case: this feature is not available to developers, but is nonetheless an essential part of the language's execution algorithm.

Other available operations are:

- Choosing a rule non-deterministically (i.e. randomly) from the available rules in a transformation.
- Executing a rule under the condition that a certain other rule does not apply.
- Executing a sequence of rules in some predefined order.
- Executing a sequence of rules in an undefined order. ATL's execution algorithm works in this way: all matched rules are executed once for each match, in an undefined order.
- Executing a rule an undefined number of times.
- Executing a rule a specified number of times.
- Repeating the execution of a rule as long as possible. After executing the rule, we check whether or not the rule still applies. If so, we execute it again.
- Trying to execute a rule, but continue if it fails to apply. One of its uses is to stop recursion over a list once all elements have been processed, and continue the execution at a higher level.

3.6. Conclusions

We analysed the structure and primitive operations within the execution algorithms of five different transformation languages. We described the commonalities and variabilities between them. In the next chapter we define, based on this analysis, our own language for defining model transformation languages in our common execution environment.

Table 3.4.: Comparison of transformation operations at the transformation level.

Primitive	GT	GROOVE	Henshin	Stratego	ATL
Execute rule	X	X	X	X	X
Match rule (no execution)				X	*
Choose a rule non-deterministically	X	X	X	X	
Execute a rule conditionally		X	X	X	
Specify a sequence of rules with defined order		X	X		
Specify a sequence of rules with undefined order					X
Execute rule an undefined number of times		X			
Execute rule a specified number of times			X		
Repeat execution of a rule as often as it is possible	X	X	X	X	
Continue if rule fails to apply		X		X	

4

Primitive Model Transformation Language

Using our observations discussed in the previous chapter, we can now define our own model transformation language, consisting of just a few primitive operations, in which all of the afore discussed transformation execution algorithms can be expressed.

4.1. Language Definition

4.1.1. Model Navigation Operations

First of all we need basic model navigation operations. The navigation operations available are quite dependent on the representation of the model in an actual implementation. Because the availability of such operations is not related to the way a transformation is executed, we assume that all operations necessary to navigate a model efficiently are available. Some example model navigation operations that are used in our primitive representations of the execution algorithms are described below.

transformation.getRules() Retrieve the rules in a transformation (graph transformation).

transformation.getSourceModel() Retrieve the source model of a transformation (graph transformation).

transformation.getTransformationUnits() Retrieve the transformation units in a transformation (Henshin).

transformation.getRules() Retrieve the rules in a transformation (Stratego).

rule.getTransformation() Retrieve the transformation to which a rule belongs (graph transformation).

rule.getTarget() Retrieve the target elements of a rule (graph transformation).

rule.getPreconditions() Retrieve the preconditions for a rule (Stratego).

unit.getSubUnits() Retrieve the subunits of a transformation unit (Henshin).

tree.getChildren() Retrieve the children of an AST element (Stratego).

4.1.2. Collection Operations

Second, we assume that basic operations to navigate and manipulate collections are also available. Some example collection operations that we use in our primitive representations of the execution algorithms are described below.

collection.add(element) Add an element as the last element of a collection.

collection.unshift() Add an element as the first element of a collection.

collection.shift() Remove the first element from a collection and return the element.

collection.remove(element) Remove an element from a collection.

collection.get(index) Retrieve the element at the specified index of a collection.

collection.contains(element) Check if a collection contains an element. Returns a boolean.

collection.copy() Returns a deep copy of a collection.

4.1.3. Logical and Arithmetic Operations

Third, we assume support for basic logical and arithmetic operations. Some examples of these operations are:

==	Equals	+	Addition
!=	Not equal to	-	Subtraction
>	Greater than	*	Multiplication
<	Less than	/	Division
=	Assignment		

4.1.4. Functions

Fourth, we assume that functions are supported. Strictly speaking, functions are not primitives, but we use them to structure the algorithms and to enable recursion. Functions can have parameters that pass references. A function does not return any value. A function named `myFunction` with a parameter named `parameter1` and type `TypeX` can be specified as follows:

```
function myFunction(TypeX parameter1):
    operation
```

4.1.5. Exception Handling

Fifth, we have exception handling. This was not observed in any of the selected transformation algorithms, but can be useful to convey information about errors during the transformation execution. An exception is raised using `raise()`. The `raise` primitive is never applicable and thus causes the execution to return to a higher level. The exception can be caught using `catch()`, which is applicable if the specified exception has been raised and was not yet caught.

4.1.6. Transformation-level Operations

The primitive transformation operations are summarized in Table 4.1. We discuss them one by one.

Find match for rule `match_rule(rule[, model])` returns Set of matches

Searches the source model for elements that match the left-hand side of the specified rule. It returns a set of matches. If the model parameter is provided, the specified (sub)model is searched instead of the source model.

We do not define the search algorithm that is used to search the model. The choice of search algorithm is left to the implementers of the primitive language.

Non-deterministic choice `nd_choice(set of expressions)` returns expression

Chooses one element of the set of expressions at random and returns the element. An expression can be an operation, a call to a function or a literal value (eg. 42).

Ordered sequence `expression ; expression`

Executes the expressions in the defined order. If the first expression fails to execute, the second will not be executed.

Continue on failure `try { expression }`

Tries to execute the expression. If it fails, continue as if the execution was successful.

Conditional execution `try { expression } else { expression }`

Tries to execute the first expression. If it fails, try to execute the second expression.

4.1.7. Rule-level Operations

Create element `model.add(element)` `create_element(class)`

The first operation (`add`) adds an element to a model. This is used in algorithms that make copies of target elements and include the copies in the target model.

The second operation (`create_element`) creates a new element that is an instance of the specified class.

Delete element `model.remove(element)`

Deletes an element from a model.

Set property `element.setProperty(name, value)`

Sets the value of the property 'name' on the element to 'value'.

Delete property `element.deleteProperty(name)`

Deletes (or nullifies) the property 'name' on the element.

Create trace `add_trace(rule, sourceElement, targetElement)`

Creates a trace that connects the source element to the target element in the context of a rule.

Query trace `match_trace_source(sourceElement)` returns Trace

Looks for a trace with the specified source element and returns it. The execution of this operation fails if no trace is found. A Trace object contains a reference to a rule, a source element and a target element.

4.1.8. Operations not included

Some of the operations observed earlier are not included in our primitive language:

- Execute a rule. This is a primitive operation at the transformation level, but one that is tied closely to the execution algorithm that is used. For example, in graph transformation, rules are executed on a per-rule basis, while in ATL the execution is split into three phases (match, create, initialize) and all rules have to be processed before the next phase is started. Therefore we cannot provide a single language construct that can be used in all algorithms.
- Execute rule an undefined number of times. This operation is only present in GROOVE and was most likely added to the language because of its roots in defining automata. We see no practical use for this operation in a general-purpose transformation language and it can be emulated with the other primitives anyway.
- Execute a rule a specified number of times. This is another operation that can be emulated using the other primitives on a per-algorithm basis, for example:

```
function for(Rule r, int n):
  try { n > 0 ; executeRule(r) ; for(r, n-1) } else { n == 0 }
```

In this example `executeRule()` is the function the algorithm in question uses to execute a transformation rule.

- Repeat execution of a rule as often as it is possible. This can be achieved by means of recursion, for example:

Table 4.1.: Primitive transformation language operations

Operation	Representation
Function	<code>function functionName(TypeX parameter1):</code>
Find match for rule	<code>match_rule(rule [, model])</code> returns Set of matches
Non-deterministic choice	<code>nd_choice(set of expressions)</code> returns expression
Ordered sequence	<code>expression ; expression</code>
Continue on failure	<code>try { expression }</code>
Conditional execution	<code>try { expression } else { expression }</code>
Create element	<code>model.add(element)</code> <code>create_element(class)</code>
Delete element	<code>model.remove(element)</code>
Set property	<code>element.setProperty(name, value)</code>
Delete property	<code>element.deleteProperty(name)</code>
Create trace	<code>add_trace(rule, sourceElement, targetElement)</code>
Query trace	<code>match_trace_source(sourceElement)</code> returns Trace
Raise exception	<code>raise(exceptionName)</code>
Catch exception	<code>catch(exceptionName)</code>

```
function repeat(Rule r): try { executeRule(r) ; repeat(r) }
```

- Specify a sequence of rules with undefined order. This can be emulated by choosing rules non-deterministically and removing them from a set (or from the transformation model) once they have been executed.

4.2. Application of the Primitive Language

We have defined our primitive model transformation language for specifying transformation algorithms, and we now express all of the aforesaid algorithms using our own language.

4.2.1. Graph Transformation

An implementation of the single-pushout graph transformation algorithm in our primitive language is included in Appendix A.1. We compare this implementation (*PL*) to the pseudo-code version (*PC*) described in Algorithm 3.1 in Chapter 3.

First, we select an arbitrary transformation rule (PL line 11, PC line 8). We search for all occurrences of the rule's left-hand side in the source model (PL line 12, PC line 10) and select an arbitrary match (PL line 14, PC line 10).

We execute the rule with this match (PL lines 22-30, PC lines 11-25). We remove all matched source elements that do not exist in the rule's right-hand side (PL line 26, PC lines 12-14). We add the elements that exist in rule's right-hand side but not in the rule's left-hand side (PL line 28, PC lines 21-23). Finally, we restart the transformation (PL line 5, PC lines 25-29).

If no match is found for a rule, we try matching another rule (PL lines 15-19, PC lines 8-10). When no rule matches, the transformation is complete (PL lines 3-6, PC lines 31-33).

4.2.2. Henshin

An implementation of Henshin's execution algorithm in our primitive language is included in Appendix A.2. This implementation reuses the `executeRule()` function from the graph transformation algorithm implementation. Again we compare this implementation (*PL*) to the pseudo-code version (*PC*) described in Algorithm 3.2 in Chapter 3.

We process all transformation units that are subunits of the Transformation one at a time (PL lines 6-8, PC lines 3-4).

- If the unit is a Rule, we look for a match in the source model and execute the Rule if one is found (PL lines 12-15, PC lines 8-11).
- If the unit is an IndependentUnit, we execute its subunits using the graph transformation algorithm (PL lines 17-9, 46-63; PC lines 13-15).
- If the unit is a PriorityUnit, we execute its subunits according to their priorities (PL lines 21-23, 65-73; PC lines 17-22).
- If the unit is a SequentialUnit, we execute its subunits in the predefined sequence (PL lines 25-27, 75-82; PC lines 24-29).
- If the unit is a CountedUnit, we execute its subunits the specified number of times (PL lines 30-31, 84-91; PC lines 31-36).
- If the unit is an AmalgamationUnit, we first try to find a match for the kernel rule (PL lines 33-35, 93-94; PL lines 38-41). If a match is found, we try to find matches for the multi-rules and execute them (PL lines 95-101, PC lines 42-44).
- If the unit is a ConditionalUnit, we first try to execute the `if` subunit (PL lines 37-38, 103-105; PC lines 46-47). If the execution of `if` succeeds, we execute the `then` subunit next (PL line 106, PC lines 48-49). If the execution of `if` fails, we execute the `else` subunit next (PL lines 107-109, PC lines 50-51).

4.2.3. Term Rewriting

An implementation of Stratego's cascading transformation strategy in our primitive language (PL) is included in Appendix A.3. The pseudo-code version (PC) was described in Algorithm 3.4 in Chapter 3.

To start, we traverse the AST to get to the bottom level (PL lines 4-6, PC lines 2-5). We try to execute the first rule in the list (PL lines 7-13, PC lines 8-9) and continue with the next rule if the first fails to execute (PL lines 16-20, PC lines 8-9). Once a successful rule execution has taken place, we restart the transformation on the current subtree (PL lines 14-15, PC lines 12-13). When no executable rules remain we move to the next level of the tree (PL lines 6-8, PC lines 4-5).

The execution of a rule was described in Algorithm 3.3 in Chapter 3. We look for matches of the left-hand side of the rule in the source model (PL line 25, PC line 3). We bind its variables to the source elements that matched (PL line 26, PC line 3). If appropriate, we check if the rule's preconditions match and bind additional variables to source elements (PL lines 27-32, PC lines 5-8). Finally, we replace the bound source elements by their respective target elements from the right-hand side of the rule (PL line 33, PC lines 9-11).

4.2.4. Declarative ATL

An implementation of the declarative ATL algorithm in our primitive language (PL) is included in Appendix A.4. The pseudo-code version (PC) was described in Algorithm 3.5 in Chapter 3.

The execution consists of three phases. First, we match the left-hand side of all rules to elements in the source model and save the matches using traces (PL lines 2, 6-15, 37-45; PC lines 7-10). Second, we create the right-hand side elements of the rule in the target model (PL lines 3, 17-25, 47-57; PC lines 12-19). Third, we initialize the properties of the target elements we just created (PL lines 4, 27-35, 59-94; PC lines 22-40).

4.2.5. Imperative ATL

We use the imperative ATL algorithm to validate that our primitive pseudo language can indeed express other kinds of transformation languages that were not considered while defining our language.

Imperative ATL uses *called rules* instead of matched rules. Called rules are not executed automatically, they must be called explicitly from another rule. Instead of defining the left-hand side and right-hand side of a rule, one defines a right-hand side and a 'do section'. The do section can contain imperative statements, namely:

- A call to a another called rule or lazy (unique) rule. Parameters can be passed to the called rule.
- An assignment statement `target <- expression` that assigns the result of an expression to a transformation module property `target`.

- A conditional statement: `if(condition) { statements1 } else { statements2 }`
- An iteration statement that iterates over a collection, executing statements one time for each element of the collection assigned to `iterator`:

```
for(iterator in collection) { statements }
```

An implementation of the imperative ATL algorithm in our primitive language is included in Appendix A.5. An imperative ATL transformation starts at the *entry point rule* (line 2). When a called rule is executed, we first create the target elements and initialize their properties (lines 6-8). Second, all statements in the *do* section are executed (line 10-55). If another rule is called, we execute this rule with the passed parameters (lines 17-18). For assignment statements we evaluate the expression and assign its result to the specified module property (lines 22-23).

For conditional statements we first evaluate the condition (lines 27-29). If the condition is true, we execute the first group of statements (line 31). If not, we execute the second group of statements (lines 32-34).

For iteration statements we iterate over the provided collection, adding the first element of the collection to the execution context of the enclosed statements each time (lines 37-55).

It should be noted that our implementation reuses some of the functions of our declarative ATL implementation. Additionally, the official ATL implementation does not create traces when executing imperative ATL. Our implementation does create traces, because it enables us to reuse part of the declarative algorithm. We consider this an implementation detail.

4.3. Related Work

Syriani and Vangheluwe [23] have conducted a similar deconstruction of transformation languages into primitives. Their primitives are objects (for instance a *Matcher*, *Rewriter*, *Iterator*) that perform parts of the transformation and exchange messages. A rough comparison between their and our primitives can be found in Table 4.2. Unfortunately, they do not motivate their choice of primitives.

Syriani and Vangheluwe do include support for parallel processing, an aspect of model transformation that we do not deal with in this thesis.

Our language includes primitives for creating and querying traces, while they do not provide support for tracing in their language.

4.4. Conclusions

We have defined our own model transformation language consisting of just a few primitive model transformation operations. We specified the transformation execution algorithms of five transformation languages in our own language, including a language

Table 4.2.: Comparison of transformation primitives

Syriani and Vangheluwe [23]	Our primitives
Matcher	Match rule
Rewriter	Create/delete element, set/delete property
Iterator	Conditional execution, functions
Resolver	Conditional execution, exception handling
Rollbacker	Continue on failure
Selector	Non-deterministic choice
Synchronizer	
Composer	Ordered sequence
	Create/query trace

with a transformation approach not initially considered. In the next chapter we work towards a transformation framework that can actually execute these specifications.

5

Transforming Model Transformations

Having defined our own primitive transformation language in Chapter 4, we can build a transformation framework that can actually transform models using transformation execution specifications written in this language.

5.1. Architecture

As mentioned before, model transformations can be described using the pattern shown in Figure 5.1. We have a source model M_a , a target model M_b and a transformation definition T_{ab} . The transformation definition is written in a transformation language TL . When executed, the transformation engine runs T_{ab} , which transforms M_a to M_b .

Our transformation engine architecture is shown in Figure 5.2. Instead of interpreting T_{ab} directly, we use it as input model for transformation T_{pl} . T_{pl} is a transformation written in our primitive language PL . It contains rules for interpreting TL transformations.

T_{pl} itself is interpreted by our primitive language interpreter I_{pl} .

Since the transformation language is itself implemented as a transformation, it is relatively easy to add support for new languages and extend existing languages.

5.2. Interpreter Implementation

The transformation engine is implemented as a plug-in for the Eclipse platform [24]. This gives us access to the Eclipse Modeling Framework [5] and related libraries.

Though the syntax of PL as described in Chapter 4 works fine for defining transformation languages, it became clear switching to a more established syntax would ease the implementation of the PL interpreter significantly. Implementing the Chapter

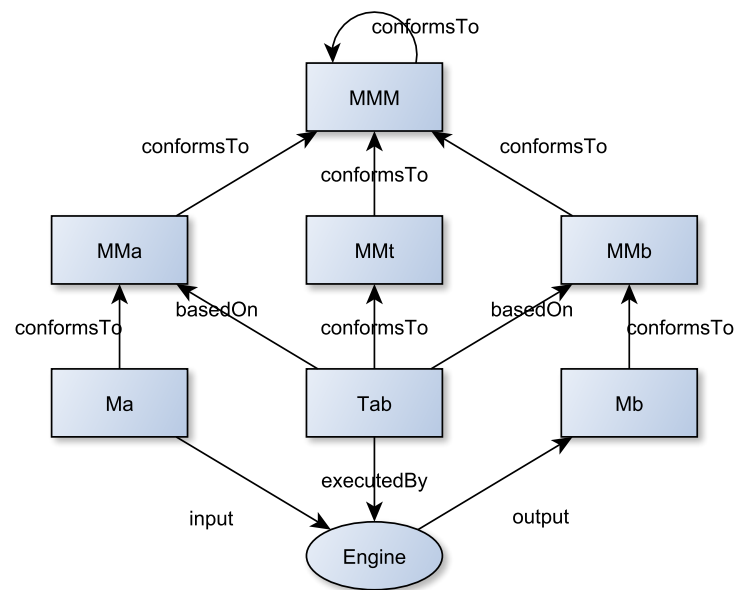


Figure 5.1.: Model Transformation Pattern




Figure 5.2.: Transformation engine architecture




Figure 5.3.: Thrascias Trace structure

4 syntax would mean implementing our own OCL [8] engine, which is required for model navigation and collection operations. This is not a small task and out of scope for this thesis.

We decided to use a syntax based on *Mistral* [9], a general purpose model transformation language and to use its implementation of OCL. We use only a small subset of Mistral's syntax, which is much more elaborate. Our implementation is called *Thrascias*.

Thrascias' abstract syntax (MM_{pl}) is defined as an ECore model, shown in Figure 5.4.

The concrete syntax is described in Table 5.1. It comprises OCL for expressions, operations for defining which models are used, a function definition and the primitive operations described in Chapter 4. In addition to the input and output models, two other metamodels are always present: the EMF metamodel Ecore and Thrascias. The Thrascias metamodel contains the Trace class which the interpreter uses for managing traces. The structure of this class is shown in Figure 5.3.

For reflective purposes, the model of the transformation currently being executed is accessible as `_Transformation_`.

We map the concrete syntax to the abstract syntax using EMFText [25]. It also provides an editor and syntax checker for the language. The interpreter itself is implemented as a standard interpreter pattern with a stack.

5.3. Transformation Language Implementation

We implement two of the analysed transformation languages in our framework: graph transformation and ATL. The implementation of the other languages is left as further work. The implementations are included in Appendix B.1 and Appendix B.2.

When implementing graph transformation there are lots of different graph transformation languages to choose from, all with different syntax and slightly different semantics. We select SimpleGT [26], because it integrates well with the Eclipse Modeling Framework and its semantics are similar to the graph transformation algorithm we analysed.




Figure 5.4.: Thrascias abstract syntax

Table 5.1.: Thrascias operations

Operation	Representation
Define transformation	<code>transform transformationName</code>
Specify input model	<code>inputModel modelName: metamodelName</code>
Specify output model	<code>outputModel modelName: metamodelName</code>
Function	<code>functionName ModelElementRule { source [parameter1: TypeX] target [result1: TypeY = expression] }</code>
Model navigation	OCL and Ecore
Collection operations	OCL
Find match for element in set	<code>match(element, elementSet) returns Sequence(Set(OclAny))</code>
Non-deterministic choice	<code>collection->any(true)</code>
Ordered sequence	<code>expression and expression</code>
Conditional execution	<code>expression or expression if expr then expr else expr endif</code>
Continue on failure	Using ModelElementRule
Create element	<code>create_element(eClass)</code>
Delete element	<code>delete element</code>
Set property	<code>update element {property=value}</code>
Delete property	<code>update element {delete property}</code>
Create trace	<code>add_trace(rule, [sourceElements[, targetElements]])</code>
Query trace	<code>match_trace_source(sourceElement) returns Trace</code>
Raise exception	<code>raise(exception)</code>
Catch exception	<code>catch(exception)</code>

5.4. Conclusions

We have explained the architecture and syntax of our transformation framework, and the rationale behind them. We wrote implementations for two transformation languages using our framework. In the next chapter we will validate our implementations with a number of common model transformation scenarios.

6

Model Transformation Scenarios

We validate our implementation by executing two different model transformation scenarios, namely transforming object-oriented classes to relation tables and pulling up subclass attributes to a superclass. First, we implement the model transformations using SimpleGT and declarative ATL. Second, we execute these implementations with both their original interpreter and our own transformation framework. Third, we compare the results of the executions.

6.1. Common Use Cases

The use cases in which model transformations are most commonly applied can be categorized into five scenarios [11]:

- Generating lower-level models and source code from higher-level models
- Reverse engineering higher-level models from lower-level models or source code
- Generating views from a model using a query
- Mapping and synchronizing models
- Model evolution or refactoring

The first three scenarios usually involve generating and parsing source code files or other documents without an explicit metamodel. Because our research focuses on model-to-model transformations, we validate our implementation using the last two scenarios.

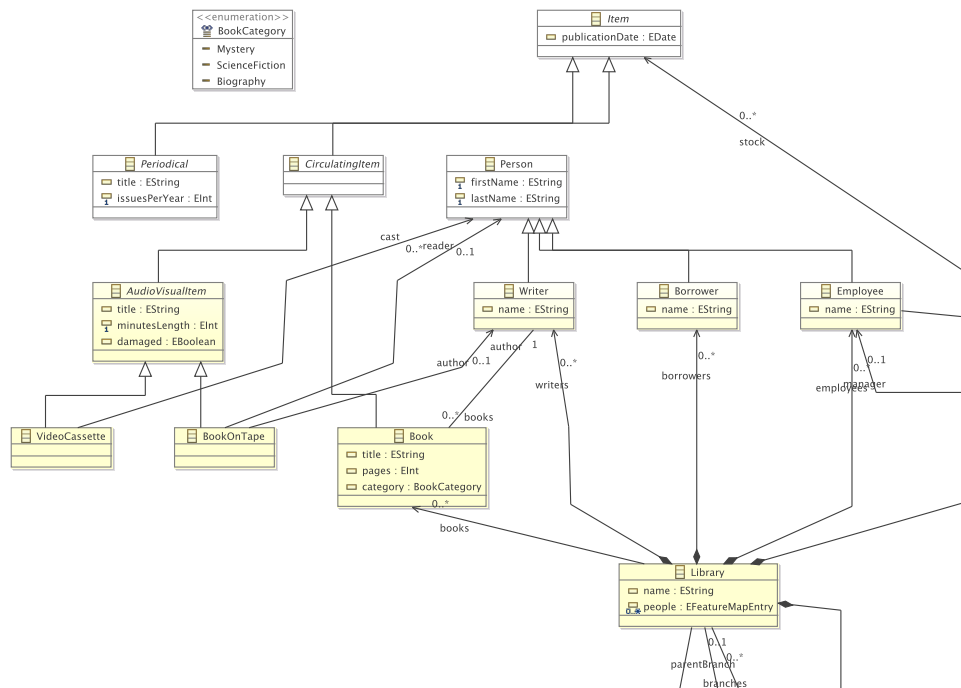


Figure 6.2.: EMF 'extlibrary' metamodel (adapted from [5])

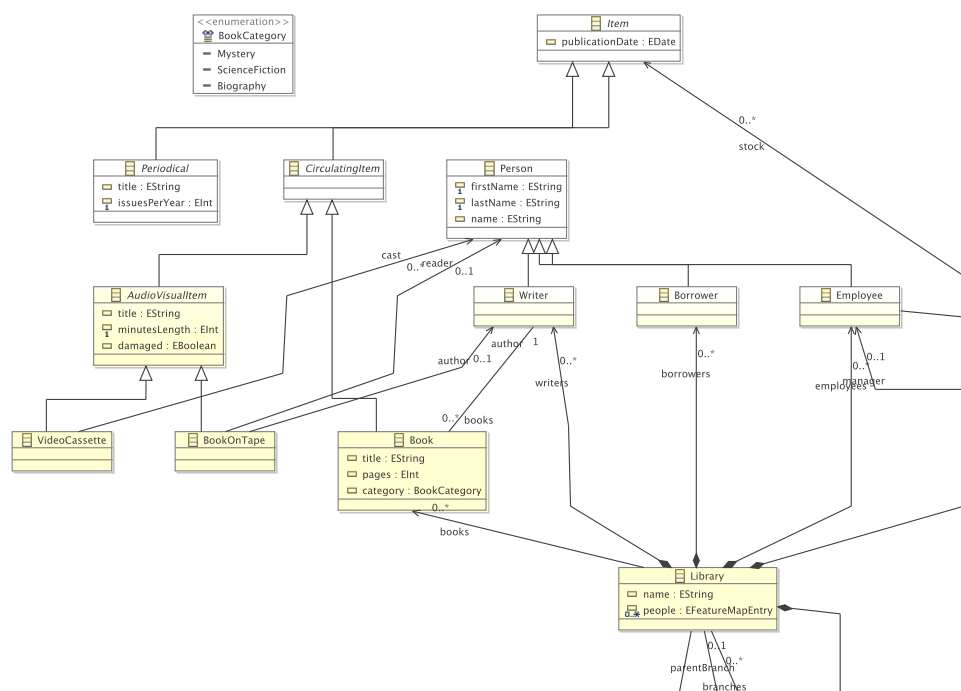


Figure 6.3.: 'extlibrary' model after Pull Up Class Attribute transformation

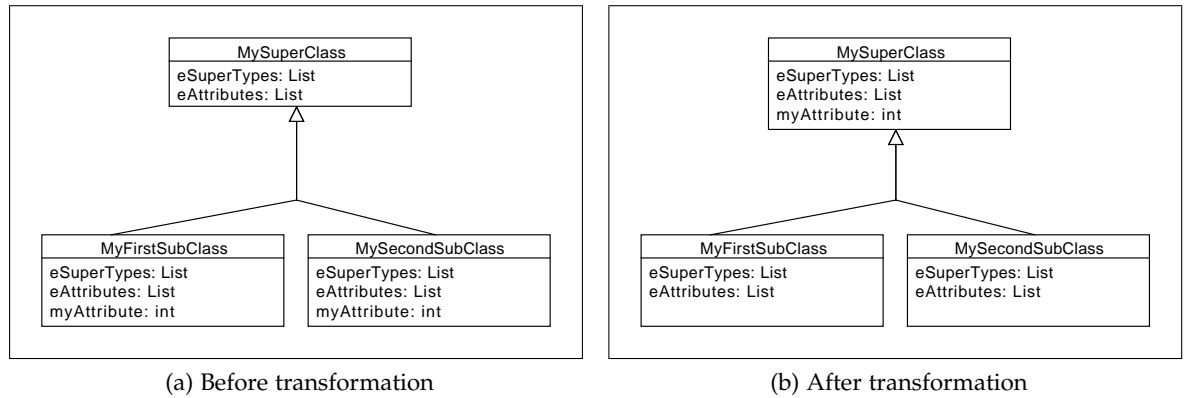


Figure 6.4.: Model for refactoring scenario

6.3. Pull Up Class Attribute to Superclass

For the second use case, model refactoring, we choose to refactor an object oriented class hierarchy. We take an attribute that is present on all subclasses of a certain class and move the attribute to the superclass. This is also known as ‘pulling up’ a class attribute.

6.3.1. Implementation

We move an attribute called `myAttribute` from a class to its superclass (Figure 6.4). The model conforms to the Ecore metamodel. Prerequisite for this refactoring is an attribute which is present in all subclasses of a certain superclass, and is not present in the superclass itself.

The implementations of this transformation in SimpleGT and declarative ATL are included in Appendix D.

6.3.2. Results

Again we apply the transformations to our example book library model (Figure 6.2). We move the attributes `name` of all the subclasses of `Person` to their superclass. The result of the transformation is shown in Figure 6.3. Verifying the results manually, we find no differences in transformation result between the original interpreters of SimpleGT or ATL and our own implementations.

6.4. Conclusions

We have introduced a number of common model transformation scenarios. We implemented two of the scenarios in SimpleGT and ATL. We then ran the transformations on

an example model using both the original interpreters and our own implementations. In our manual verification we found no difference in transformation result.

In the next chapter we summarize the results of our research and look at possible future directions of research.

7

Conclusions and Future Work

In Chapter 1 we introduced the problem of having many different model transformation languages, all with their own unique approach. We recognized that, despite their differences, they have some features in common. We asked the question: “how can we execute different model transformation languages in a common environment, while retaining the variability between the languages?”. We can now answer this question based on our research.

7.1. What should a common execution environment for model transformation languages look like?

In Chapter 2 we explored models and model transformations. We categorized transformation languages based on properties like transformation approach and rule application strategy. Using this categorization, we selected five significantly different transformation languages for further investigation: graph transformation, GROOVE, Henshin, Stratego and declarative ATL.

In Chapter 3 we took a closer look at the execution algorithms of the selected languages. We analysed the commonality and variability between them, both at the rule level and the transformation level.

In Chapter 4 we defined our own primitive model transformation language based on the analysis of commonality and variability. We expressed the execution algorithms for all analysed languages in our primitive language.

Subsequently, we took a look at imperative ATL. It has a completely different transformation approach, one not considered during the design of our primitive language. Nevertheless we were able to express its execution algorithm, demonstrating

that our language can be used to describe the execution algorithm of any transformation language.

7.2. How do we build support for executing different transformation languages in this common environment?

In Chapter 5 we discussed the implementation of a model transformation framework based on our primitive transformation language. We described the architecture, syntax and interpreter design. We implemented two model transformation languages in our framework: SimpleGT and declarative ATL.

In Chapter 6 we tested our transformation framework by executing two common model transformations in both our own interpreter and the original interpreter for the transformation language. For our humble test scenarios, the result turned out to be the same regardless of the interpreter used.

7.3. What are the limitations of executing transformation languages in a common environment?

We have shown that, despite the variability between them, transformation languages have lots of commonalities. Their execution algorithms can be described using less than fifteen primitive operations. Based on these findings we have implemented our own transformation framework for executing model transformation written in these languages. However, our implementation has some limitations.

Model representation To ease the implementation of our framework, we require all models to be represented as EMF Ecore models. In some of the discussed transformation languages models are usually represented differently. Future work could include creating an abstraction layer which maps these model representations to Ecore.

OCL for navigation and collection manipulation Another generalization is the use of OCL for model navigation in all languages. Other navigation languages can be used, but one would have to add an interpreter for such a language either to the Thrascias interpreter or to the transformation language implementation.

Multiple source and target models To prevent the example implementations from becoming overly complex we did not consider transforming from multiple source models, or transforming to multiple target models. This however does not mean there is any theoretical or practical impediment to implement this.

Limited number of implemented languages We implemented just two of the five languages analysed. Furthermore there are many more transformation languages available. Future work could include the implementation of the other three languages mentioned in this thesis, or any other available transformation language.

Limited number of test scenarios We performed tests using two relatively simple transformation scenarios. Testing larger and more advanced transformations can be considered future work.

7.4. Looking ahead

Besides addressing some of the limitations of our implementation, we take a moment to look at possible directions for future research.

We have shown that different transformation languages can be interpreted with a single interpreter. A logical next step would be the composition of transformations written in different transformation languages into a single transformation. This would enable one to create libraries of transformations for common transformation tasks that can be used in no matter which language. Moreover a separate transformation language could be chosen for each transformation subproblem. Combining transformation languages introduces a host of new problems, like different composition strategies and conflict resolution between rules in different languages.

A second direction for future research could be parallel processing. Our framework does not take advantage of multiple processing units at the moment. Additional language constructs might be necessary for operations like synchronization of parallel threads.

Performance improvement is a possible third direction of future research. Although we have not conducted any testing for processing speed or memory usage, we expect our prototype implementation to be slower than the original interpreters for most languages.



Primitive Representations of the Algorithms

A.1. Single-Pushout Graph Transformation

```
1 function executeTransformation(Transformation transformation):
2     // repeat as often as possible
3     try {
4         transformGraph(transformation.getRules()) ;
5         executeTransformation(transformation)
6     }
7
8 function transformGraph(Set rules):
9     try {
10         // choose random rule and find match
11         Rule rule = nd_choice(rules) ;
12         Set matches = match_rule(rule) ;
13         // execute this rule
14         executeRule(rule, nd_choice(matches))
15     } else { // no matches
16         // remove rule from list
17         rules = rules.remove(rule) ;
18         // try again
19         transformGraph(rules)
20     }
21
22 function executeRule(Rule rule, Match match):
23     // make temporary copy of graph
24     Model D = rule.getTransformation().getSourceModel().copy() ;
25     // remove elements not in target
26     removeOldElements(D, rule, match.copy()) ;
27     // add elements not in source
28     addNewElements(D, match, rule.getTarget()) ;
29     // replace graph by changed copy
```

```
30     rule.getTransformation().setSourceModel(D)
31
32 function removeOldElements(Model D, Rule rule, Match match):
33     try {
34         ModelElement element = nd_choice(match.getSource()) ;
35         try {
36             rule.getTarget().contains(element)
37         } else {
38             D.remove(element)
39         } ;
40         removeOldElements(D, rule, match.remove(element))
41     }
42
43 function addNewElements(D, match, rhs):
44     try {
45         ModelElement element = nd_choice(rhs) ;
46         try {
47             match.getSource().contains(element)
48         } else {
49             D.add(element)
50         } ;
51         addNewElements(D, rhs.remove(element))
52     }
```

A.2. Henshin

```

1  function executeTransformation(Transformation transformation):
2      processAllUnits(transformation.getTransformationUnits())
3
4  function processAllUnits(Set units):
5      try {
6          TransformationUnit unit = nd_choice(unit) ;
7          processUnit(unit) ;
8          processAllUnits(units.remove(unit))
9      }
10
11 function processUnit(TransformationUnit unit):
12     try {
13         unit.isRule() ;
14         Set rule_matches = match_rule(unit) ;
15         executeRule(unit, nd_choice(matches))
16     } else {
17         try {
18             unit.isIndependentUnit() ;
19             processIndependentUnit(unit)
20         } else {
21             try {
22                 unit.isPriorityUnit() ;
23                 processPriorityUnit(unit, null)
24             } else {
25                 try {
26                     unit.isSequentialUnit() ;
27                     processSequentialUnit(unit)
28                 } else {
29                     try {
30                         unit.isCountedUnit() ;
31                         processCountedUnit(unit, unit.getCount())
32                     } else {
33                         try {
34                             unit.isAmalgamationUnit() ;
35                             processAmalgamationUnit(unit)
36                         } else {
37                             unit.isConditionalUnit() ;
38                             processConditionalUnit(unit)
39                         }
40                     }
41                 }
42             }
43         }
44     }
45
46 function processIndependentUnit(IndependentUnit unit):
47     // repeat as often as possible
48     try {
49         processIndependentSubUnits(unit.getSubUnits()) ;
50         processIndependentUnit(unit)
51     }

```

```

52
53 function processIndependentSubUnits(Set units):
54     try {
55         // process random unit
56         TransformationUnit unit = nd_choice(units) ;
57         processUnit(unit)
58     } else {
59         // remove unit from list
60         units = units.remove(unit) ;
61         // try again
62         processIndependentSubUnits(units)
63     }
64
65 function processPriorityUnit(PriorityUnit unit, TransformationUnit
    currentSubUnit):
66     try {
67         TransformationUnit nextSubUnit = unit.getNextSubUnit(currentSubUnit)
68         ;
69         try {
70             processUnit(nextSubUnit)
71         } else {
72             processPriorityUnit(unit, nextSubUnit)
73         }
74     }
75
76 function processSequentialUnit(SequentialUnit unit, TransformationUnit
    currentSubUnit):
77     try {
78         TransformationUnit nextSubUnit = unit.getNextSubUnit(currentSubUnit)
79         ;
80         processUnit(nextSubUnit) ;
81         processSequentialUnit(unit, nextSubUnit)
82     } else {
83         nextSubUnit == null
84     }
85
86 function processCountedUnit(CountedUnit unit, Integer counter):
87     try {
88         counter > 0 ;
89         processUnit(unit.getSubUnit()) ;
90         processCountedUnit(unit, counter - 1)
91     } else {
92         counter == 0
93     }
94
95 function processAmalgamationUnit(AmalgamationUnit unit):
96     processUnit(unit.getKernelRule()) ;
97     processMultiRules(unit.getMultiRules())
98
99 function processMultiRules(Set rules):
100     try {
101         processIndependentSubUnits(rules) ;
102         processMultiRules(rules)
103     }

```



```
102
103 function processConditionalUnit(ConditionalUnit unit):
104     try {
105         processUnit(unit.getIf()) ;
106         processUnit(unit.getThen())
107     } else {
108         processUnit(unit.getElse())
109     }
```

A.3. Stratego

```

1  function executeTransformation(Transformation transformation):
2      cascade(transformation.getRules(), transformation.getSourceModel())
3
4  function cascade(List rules, AST tree):
5      try {
6          cascade(rules, nd_choice(tree.getChildren()))
7      } ;
8      applyAndRestart(rules, tree, rules)
9
10 function applyAndRestart(List rules, AST tree, List allRules):
11     Rule rule = rules.shift() ;
12     try {
13         applyRule(rule, tree) ;
14         cascade(allRules, tree) ;
15         raise(BreakException)
16     } else {
17         try {
18             catch(BreakException)
19         } else {
20             applyAndRestart(rules, tree, allRules)
21         }
22     }
23
24 function applyRule(Rule rule, AST tree):
25     Match m = nd_choice(match_rule(rule, tree)) ;
26     Set bindings = m.getBindings() ;
27     try {
28         Match s = match_rule(rule.getPreconditions(), tree) ;
29         bindings.append(s.getBindings())
30     } else {
31         rule.getPreconditions() == null
32     } ;
33     tree = rule.getTarget().replaceVars(bindings)

```

A.4. Declarative ATL

```

1  function executeTransformation(Transformation transformation):
2      matchAllRules(transformation.getRules()) ;
3      createElements(transformation.getAllTraces()) ;
4      initElements(transformation.getAllTraces())
5
6  function matchAllRules(Set rules):
7      try {
8          // choose rule randomly
9          Rule rule = nd_choice(rules) ;
10         // find matches and save them
11         Set rule_matches = match_rule(rule) ;
12         findMatches(rule, rule_matches) ;
13         // match next rule
14         matchAllRules(rules.remove(rule))
15     }
16
17  function createElements(Set traces):
18      try {
19          // choose trace (= match) randomly
20          Trace trace = nd_choice(traces) ;
21          // create targets for trace
22          createTargets(trace, trace.getRule().getTarget().copy()) ;
23          // proceed with next trace
24          createElements(traces.remove(trace))
25      }
26
27  function initElements(Set traces):
28      try {
29          // choose trace randomly
30          Trace trace = nd_choice(traces) ;
31          // init its targets
32          initTargets(trace.getTarget().copy()) ;
33          // proceed with next trace
34          initElements(targets.remove(targetElement))
35      }
36
37  function findMatches(Rule rule, Set matches):
38      try {
39          // get match
40          Match match = nd_choice(matches)
41          // add trace with empty target
42          add_trace(rule, match.getSource(), null) ;
43          // process next match
44          findMatches(rule, matches.remove(match))
45      }
46
47  function createTargets(Trace trace, Set targets):
48      try {
49          // choose target element randomly
50          ModelElement targetElement = nd_choice(targets) ;
51          // create empty target element

```

```

52         Object t = trace.getRule().getTransformation().getTargetModel()
53             .add(create_element(targetElement.class)) ;
54         // set target element for trace
55         trace.getTarget().add(t) ;
56         createTargets(trace, targets.remove(targetElement))
57     }
58
59     function initTargets(Set targets):
60         try {
61             ModelElement targetElement = nd_choice(targets) ;
62             initProperties(targetElement, targetElement.getProperties().copy()) ;
63             initTargets(targets.remove(targetElement))
64         }
65
66     function initProperties(ModelElement target, Set properties):
67         try {
68             // choose property randomly
69             ModelElementProperty property = nd_choice(properties) ;
70             try {
71                 // if primitive, assign directly
72                 property.getContent().isPrimitive() ;
73                 target.setProperty(property.getName(), property.getContent())
74             } else {
75                 try {
76                     // if source element, resolve, then assign
77                     Trace trace = match_trace_source(property.getContent()) ;
78                     target.setProperty(property.getName(),
79                                     trace.getTarget().get(0))
80                 } else {
81                     try {
82                         // if trace query, resolve, then assign
83                         property.getContent().isQuery() ;
84                         Trace trace = match_trace_source(property.getContent()
85                                                         .getSource()) ;
86                         target.setProperty(property.getName(), trace.getTarget()
87                                                         .get(property.getContent().getPattern()))
88                     } else {
89                         // if target element, assign directly
90                         target.setProperty(property.getName(),
91                                         property.getContent())
92                     }
93                 }
94             } ;
95             initProperties(target, properties.remove(property))
96         }

```

A.5. Imperative ATL

```

1 function executeTransformation(Transformation transformation):
2     executeRule(transformation.getEntryPointRule(), null)
3
4 function executeRule(Rule rule, List parameters):
5     // Execute 'to' block
6     Trace trace = add_trace(rule, null, null) ;
7     createTargets(trace, rule.getTarget().copy()) ;
8     initTargets(trace.getTarget().copy()) ;
9     // Execute 'do' block
10    executeStatements(rule, parameters, rule.getStatements())
11
12 function executeStatements(Rule rule, List parameters, List statements):
13     try {
14         Statement s = statements.get(0) ;
15         try {
16             // Call another rule
17             s.isCallStatement() ;
18             executeRule(s.getRule(), s.getParameters())
19         } else {
20             try {
21                 // Assign module property
22                 s.isAssignmentStatement() ;
23                 rule.getModule().setProperty(s.getTarget(),
24                     s.getExpression())
25             } else {
26                 try {
27                     // Execute 'if' statement
28                     s.isIfStatement() ;
29                     try {
30                         s.getCondition() ;
31                         // Condition = true
32                         executeStatements(rule, parameters,
33                             s.getStatements1())
34                     } else {
35                         // Condition = false
36                         executeStatements(rule, parameters,
37                             s.getStatements2())
38                     }
39                 } else {
40                     try {
41                         // Execute 'for' statement
42                         s.isForStatement() ;
43                         executeForStatement(rule, parameters,
44                             s.getCollection(),
45                             s.getStatements())
46                     }
47                 }
48             }
49         } ;
50     } ;
51     executeStatements(rule, parameters, statements.unshift())
52 }

```

```
48
49 function executeForStatement(Rule rule, List parameters, Collection
    collection,
50     List statements):
51     try {
52         List newParameters = parameters.copy().add(collection.get(0)) ;
53         executeStatements(rule, newParameters, statements) ;
54         executeForStatement(rule, parameters, collection.unshift(),
            statements)
55     }
```

B

Thrascias Implementations of the Algorithms

B.1. SimpleGT

```
1 transform simplegt
2 inputModel SimpleGT : simplegt
3 outputModel Out : simplegt
4
5 executeTransformationRule ModelElementRule {
6     source [transformation: SimpleGT!Module]
7     target [
8         -- run transformation, resulting in a new model
9         transformStep: Tuple() =
10             transformGraphRule(OrderedSet{transformation.models}
11                 ->first(), transformation.elements),
12         -- replace old source model by new model
13         update transformation {models = OrderedSet{transformStep.newModel}},
14         -- if a rule was applied, we run the transformation again
15         result: SimpleGT!OclModel =
16             if not transformStep.newModel.oclIsUndefined() then
17                 executeTransformationRule(transformation)
18             else
19                 transformation.models ->first()
20             endif
21     ]
22 }
23 transformGraphRule ModelElementRule {
24     source [sourceModel: SimpleGT!OclModel, rules: OrderedSet(SimpleGT!Rule)]
25     target [
26         -- select a random rule and try to find a match
27         rule: SimpleGT!Rule = rules ->asSet()->any(true),
28         matches: Set(Set(OclAny)) =
```

```

29         if not rule.oclIsUndefined() then
30             match(rule.input.elements, sourceModel)
31         else
32             -- no rules left
33             Set{}
34         endif,
35     newModel: SimpleGT!OclModel =
36         if matches ->notEmpty() and not
37             hasApplicableNacRule(sourceModel, rule).result
38         then
39             -- match found, execute the rule for this match
40             executeRule(sourceModel, rule, matches ->any(true)).result
41         else
42             -- no match, remove this rule from the set and try again
43             transformGraphRule(rules ->excluding(rule)).newModel
44         endif
45     ]
46 }
47 hasApplicableNacRule ModelElementRule {
48     source [sourceModel: SimpleGT!OclModel, rule: SimpleGT!Rule]
49     target [
50         result: Boolean = rule.nac ->exists(n | match(n.elements,
51             sourceModel) ->notEmpty())
52     ]
53 }
54 executeRule ModelElementRule {
55     source [sourceModel: SimpleGT!OclModel, rule: SimpleGT!Rule, match:
56         Set(OclAny)]
57     target [
58         -- remove elements that are not in the target graph
59         D: OrderedSet(ETCore!EObject) = removeOldElementsRule(sourceModel,
60             rule, match)
61         .result,
62         -- check for dangling edges
63         updatedObjectSet: OrderedSet(ETCore!EObject) =
64             if hasDanglingEdgeRule(D ->select(o |
65                 o.oclIsKindOf(ETCore!EReference)))
66                 .result then
67                     -- abort rule execution
68                     sourceModel.elements
69                 else
70                     -- add elements that are not in the source graph
71                     D ->union(rule.output.elements ->collect(e |
72                         create_element(e.eClass()))
73                     ->asOrderedSet())
74                 endif,
75         -- update model
76         update sourceModel {elements = updatedObjectSet},
77         result: SimpleGT!OclModel = sourceModel
78     ]
79 }

```



```
77 removeOldElementsRule ModelElementRule {
78     source [sourceModel: SimpleGT!OclModel, rule: SimpleGT!Rule, match:
79         Set(OclAny)]
80     target [
81         oldElements: Set(OclAny) = match - rule.output.elements,
82         result: OrderedSet(ECore!EObject) = sourceModel.elements -
83             oldElements
84     ]
85 }
86
87 hasDanglingEdgeRule ModelElementRule {
88     source [edges: Collection(ECore!EReference)]
89     target [
90         result: Boolean = edges ->exists(e| e.eContainer.isOclUndefined()
91             or e.getEType().isOclUndefined())
92     ]
93 }
```

B.2. Declarative ATL

```

1  transform atl
2  inputModel ATL: atl
3  outputModel Out: atl
4
5  executeTransformationRule ModelElementRule {
6      source [transformation: ATL!Module]
7      target [
8          matches: Tuple() =
9              matchAllRulesRule(OrderedSet{transformation.inModels}->first(),
10                 OrderedSet{transformation.elements->select(e|
11                     e.ocIsKindOf(ATL!MatchedRule))},
12                 elements: Tuple() = createElementsRule(OrderedSet{matches.traces}),
13                 result: Tuple() = initElementsRule(OrderedSet{matches.traces})
14             ]
15     }
16
17     -- Find all matches for rules and save them using traces
18     matchAllRulesRule ModelElementRule {
19         source [sourceModel: ATL!OclModel, rules: OrderedSet(ATL!MatchedRule)]
20         target [
21             rule: ATL!MatchedRule = rules ->asSet()->any(true),
22             matches: Sequence(Set(OclAny)) =
23                 if not rule.ocIsUndefined() then
24                     -- match rule in source model
25                     match(rule.inPattern.elements, sourceModel)
26                 else
27                     -- no rules left to match
28                     Sequence{}
29                 endif,
30             -- create trace with empty target
31             traces: Bag(Thrascias!Trace) = matches ->collect(m| add_trace(rule,
32                 m)),
33             result: Bag(Thrascias!Trace) =
34                 if rules ->size() > 1 then
35                     -- match next rule
36                     matchAllRulesRule(sourceModel, rules
37                         ->excluding(rule))->including(matches)
38                 else
39                     Bag{}
40                 endif
41         ]
42     }
43
44     -- Create all new elements using the traces
45     createElementsRule ModelElementRule {
46         source [traces: OrderedSet(Thrascias!Trace)]
47         target [
48             result: Collection(Tuple()) = traces ->collect(tr|
49                 tr.target ->collect(ta| createTargetsRule(tr, ta,
50                     OrderedSet{tr.rule.module.outModels}->first()))
51         ]
52     }
53 }

```

```

48     ]
49 }
50
51 -- Create new element for trace in output model
52 createTargetsRule ModelElementRule {
53     source [trace: Thrascias!Trace, targetElement: Ecore!EObject, outModel:
54         ATL!OclModel]
55     target [
56         -- create empty target element
57         newElement: OclAny = create_element(targetElement.eClass()),
58         -- set target element for trace
59         update trace {target = trace.target ->including(newElement)},
60         -- add element to output model
61         update outModel {elements = outModel.elements
62             ->including(newElement)}
63     ]
64 }
65
66 -- Initialize properties for all new elements
67 initElementsRule ModelElementRule {
68     source [traces: OrderedSet(Thrascias!Trace)]
69     target [
70         result: OclAny = traces ->collect(tr|
71             tr.source ->collect(s|
72                 -- find input pattern element for trace
73                 let ipElement: ATL!InPatternElement = match(s,
74                     rule.inPattern.elements)
75                     ->any(m| m.notEmpty()) in
76                     -- find target elements and initialize their properties
77                     ipElement.mapsTo -> collect(opElement|
78                         initTargets(opElement, match(opElement, tr.target)
79                             ->any(m| m.notEmpty()))
80                 )
81             )
82         ]
83     ]
84 }
85
86 -- Initialize all properties of a target element
87 initTargets ModelElementRule {
88     source [outputElement: ATL!OutPatternElement, targetElement:
89         Ecore!EObject]
90     target [
91         result: Bag(Tuple()) = outputElement.bindings ->collect(b|
92             resolvePropertyRule(targetElement, b))
93     ]
94 }
95
96 -- Resolve the value of a property
97 resolvePropertyRule ModelElementRule {
98     source [targetElement: Ecore!EObject, property: ATL!Binding]
99     target [
100         result: OclAny = if property.value.type.ocIsKindOf(ATL!Primitive)
101             then

```

```

97         -- property has primitive value, initialize directly
98         initPropertyRule(targetElement, property, property.value)
99     else
100         -- if property has source element value, resolve, then
            initialize
101         let trace: Thrascias!Trace =
            match_trace_source(property.value) in
102         if not trace.oclIsUndefined() then
103             initPropertyRule(targetElement, property,
                trace.target)
104         else
105             -- if property has trace query value, resolve, then
                initialize
106             -- example: resolveTemp(source_element,
                target_pattern)
107             if property.value.oclIsKindOf(ATL!OperationCallExp)
                and
108             property.value.operationName = 'resolveTemp' then
109                 let trace: Thrascias!Trace =
110                     match_trace_source(property.value.arguments
                        ->at(1)) in
111                 let sourceObject: Ecore!Object =
112                     OrderedSet{trace.target}->first() in
113                     initPropertyRule(targetElement, property,
                        sourceObject
                            .eGet(sourceObject.eClass().getEAllStructuralFeatures
                                ->any(f| f.getName() =
                                    property.value.arguments
                                        ->at(2))))
114             else
115                 -- if property has target element value,
                    initialize directly
116                 initPropertyRule(property, property.value)
117             endif
118         endif
119     endif
120 ]
121 }
122
123 -- Initialize the value of a property
124 initPropertyRule ModelElementRule {
125     source [targetElement: Ecore!EObject, property: ATL!Binding, value:
        OclAny]
126     target [
127         -- find structural feature for property
128         feature: Ecore!EStructuralFeature = targetElement.eClass()
129             .getEAllStructuralFeatures() ->any(f| f.getName() =
                property.propertyName),
130         -- set value of feature
131         result: Boolean = targetElement.eSet(feature, value)
132     ]
133 }

```



Implementations of the Mapping Scenario

C.1. SimpleGT

```
1 module Class2Relational;
2
3 metamodel Class;
4 metamodel Relational;
5 transform C: Class, R: Relational;
6
7 rule Class2Table {
8   from c : Class!Class
9   to t : Relational!Table (
10     name =~ c.name,
11     col =~ Sequence {key}->union(c.attr->select(e | not
12       e.multivalued)),
13     key =~ Set{key}
14   ),
15   key : Relational!Column (
16     name =~ 'objectId',
17     type =~ Class!DataType.allInstances()->select(e | e.name =
18       'Integer')->first()
19   )
20 }
21
22 rule DataType2Type {
23   from dt : Class!DataType
24   to t : Relational!Type (name =~ dt.name)
25 }
26
27 rule DataTypeAttribute2Column {
28   from a : Class!Attribute (
29     type =~ dt,
```

```

28         multivalued =~ false
29     ),
30     dt : Class!DataType
31 }
32
33 rule MultivaluedDataTypeAttribute2Column {
34     from a : Class!Attribute (
35         type =~ dt,
36         multivalued =~ true
37     ),
38     dt : Class!DataType
39     to t : Relational!Table (
40         name =~ a.owner.name + '_' + a.name,
41         col =~ Sequence{id, value}
42     ),
43     id : Relational!Column (
44         name =~ a.owner.name + 'Id',
45         type =~ Class!DataType.allInstances()->select(e | e.name =
46             'Integer')->first()
47     ),
48     value : Relational!Column (
49         name =~ a.name,
50         type =~ a.type
51     )
52 }
53
54 rule ClassAttribute2Column {
55     from a : Class!Attribute (
56         type =~ c,
57         multivalued =~ false
58     ),
59     c : Class!Class
60     to foreignKey : Relational!Column (
61         name =~ a.name + 'Id',
62         type =~ Class!DataType.allInstances()->select(e | e.name =
63             'Integer')->first()
64     )
65 }
66
67 rule MultivaluedClassAttribute2Column {
68     from a : Class!Attribute (
69         type =~ c,
70         multivalued =~ true
71     ),
72     c : Class!Class
73     to t : Relational!Table (
74         name =~ a.owner.name + '_' + a.name,
75         col =~ Sequence{id, foreignKey}
76     ),
77     id : Relational!Column (
78         name =~ a.owner.name + 'Id',
79         type =~ Class!DataType.allInstances()->select(e | e.name =
80             'Integer')->first()
81     ),

```

```
79         foreignKey : Relational!Column (
80             name =~ a.name + 'Id',
81             type =~ Class!DataType.allInstances()->select(e | e.name =
                'Integer')->first()
82         )
83
84 }
```

C.2. Declarative ATL

```

1  -- Adapted from ATL Modeling Zoo [27]
2
3  module Class2Relational;
4  create OUT : Relational from IN : Class;
5
6  rule Class2Table {
7      from
8          c : Class!Class
9      to
10         out : Relational!Table (
11             name <- c.name,
12             col <- Sequence {key}->union(c.attr->select(e | not
13                 e.multivalued)),
14             key <- Set {key}
15         ),
16         key : Relational!Column (
17             name <- 'objectId',
18             type <- Class!DataType.allInstances()->select(e | e.name =
19                 'Integer')->first()
20         )
21     }
22
23 rule DataType2Type {
24     from
25         dt : Class!DataType
26     to
27         out : Relational!Type (
28             name <- dt.name
29         )
30     }
31
32 rule DataTypeAttribute2Column {
33     from
34         a : Class!Attribute (
35             a.type.ocIsKindOf(Class!DataType) and not a.multivalued
36         )
37     to
38         out : Relational!Column (
39             name <- a.name,
40             type <- a.type,
41             owner <- thisModule.resolveTemp(a.owner, 'key')
42         )
43     }
44
45 rule MultivaluedDataTypeAttribute2Column {
46     from
47         a : Class!Attribute (
48             a.type.ocIsKindOf(Class!DataType) and a.multivalued
49         )
50     to
51         out : Relational!Table (

```



```

50         name <- a.owner.name + '_' + a.name,
51         col <- Sequence {id, value}
52     ),
53     id : Relational!Column (
54         name <- a.owner.name + 'Id',
55         type <- Class!DataType.allInstances()->select(e | e.name =
                    'Integer')->first()
56     ),
57     value : Relational!Column (
58         name <- a.name,
59         type <- a.type
60     )
61 }
62
63 rule ClassAttribute2Column {
64     from
65         a : Class!Attribute (
66             a.type.ocIsKindOf(Class!Class) and not a.multivalued
67         )
68     to
69         foreignKey : Relational!Column (
70             name <- a.name + 'Id',
71             type <- Class!DataType.allInstances()->select(e | e.name =
                    'Integer')->first()
72         )
73 }
74
75 rule MultivaluedClassAttribute2Column {
76     from
77         a : Class!Attribute (
78             a.type.ocIsKindOf(Class!Class) and a.multivalued
79         )
80     to
81         t : Relational!Table (
82             name <- a.owner.name + '_' + a.name,
83             col <- Sequence {id, foreignKey}
84         ),
85         id : Relational!Column (
86             name <- a.owner.name + 'Id',
87             type <- Class!DataType.allInstances()->select(e | e.name =
                    'Integer')->first()
88         ),
89         foreignKey : Relational!Column (
90             name <- a.name + 'Id',
91             type <- Class!DataType.allInstances()->select(e | e.name =
                    'Integer')->first()
92         )
93 }

```


D

Implementations of the Refactoring Scenario

D.1. SimpleGT

```
1 module PullUpAttribute;
2
3 metamodel Ecore;
4 transform IN: Ecore, OUT: Ecore;
5
6 rule AddAttributeToSuperclass {
7   from sc : Ecore!EClass,
8       c : Ecore!EClass (
9       eSuperTypes =~ sc,
10      eAttributes =~ a
11    ),
12   a : Ecore!EAttribute (
13     name =~ 'myAttribute'
14   )
15   not sc : Ecore!EClass (
16     eAttributes =~ a
17   ),
18   c2 : Ecore!EClass (
19     eSuperTypes =~ sc,
20     eAttributes =~ c2.eAttributes->excluding(a)
21   )
22   to out : Ecore!EClass (
23     name =~ sc.name,
24     eSuperTypes =~ sc.eSuperTypes,
25     eAttributes =~ sc.eAttributes->including(a)
26   )
27 }
28
29 rule RemoveAttributeFromSubclass {
```

```
30     from sc : Ecore!EClass,  
31         c : Ecore!EClass (  
32             eSuperTypes =~ sc,  
33             eAttributes =~ a  
34         ),  
35         a : Ecore!EAttribute (  
36             name =~ 'myAttribute'  
37         )  
38     not sc : Ecore!EClass (  
39         eAttributes =~ a  
40     ),  
41     c2 : Ecore!EClass (  
42         eSuperTypes =~ sc,  
43         eAttributes =~ c2.eAttributes->excluding(a)  
44     )  
45     to out : Ecore!EClass (  
46         name =~ c.name,  
47         eSuperTypes =~ c.eSuperTypes,  
48         eAttributes =~ c.eAttributes->excluding(a)  
49     )  
50 }
```

D.2. Declarative ATL

```

1  module PullUpAttribute;
2  create OUT : Ecore from IN : Ecore
3
4  rule AddAttributeToSuperclass {
5      from
6          c : Ecore!EClass (
7              Ecore!EClass.allInstances()->forall(e|
8                  c.isSuperTypeOf(e) implies e.getEAttributes()->exists(a|
9                      a.name = ca.name))
10             and not c.getEAttributes()->exists(a| a.name = ca.name)
11         ),
12         ca : Ecore!EAttribute (
13             ca.name = 'myAttribute'
14         )
15     to
16         out : Ecore!EClass (
17             name <- c.name,
18             eSuperTypes <- c.eSuperTypes,
19             eAttributes <- Sequence{c.getEAttributes()}->including(ca)
20         )
21 }
22
23 rule RemoveAttributeFromSubclass {
24     from
25         c : Ecore!EClass (
26             c.eAttributes->exists(a| a.name = ca.name) and
27             c.eSuperTypes->exists(s|
28                 Ecore!EClass.allInstances()->forall(e| s.isSuperTypeOf(e)
29                     implies (
30                         e.getEAttributes()->exists(a| a.name = ca.name)
31                         and not s.getEAttributes()->exists(a| a.name = ca.name)
32                     )))
33     ),
34     ca : Ecore!EAttribute (
35         ca.name = 'myAttribute' and
36         ca.eContainer() = c
37     )
38 to
39     out : Ecore!EClass (
40         name <- c.name,
41         eSuperTypes <- c.eSuperTypes,
42         eAttributes <- Sequence{c.getEAttributes()}->excluding(ca)
43     )
44 }

```


Bibliography

- [1] E. Biermann, K. Ehrig, C. Köhler, G. Kuhns, G. Taentzer, and E. Weiss, “Graphical definition of in-place transformations in the Eclipse Modeling Framework,” in *Model Driven Engineering Languages and Systems*, ser. Lecture Notes in Computer Science, O. Nierstrasz, J. Whittle, D. Harel, and G. Reggio, Eds. Springer Berlin / Heidelberg, 2006, vol. 4199, pp. 425–439. [Online]. Available: http://dx.doi.org/10.1007/11880240_30
- [2] F. Jouault, F. Allilaire, J. Bézivin, and I. Kurtev, “ATL: A model transformation tool,” *Science of Computer Programming*, vol. 72, no. 1–2, pp. 31 – 39, 2008. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167642308000439>
- [3] T. Arendt, E. Biermann, S. Jurack, C. Krause, and G. Taentzer, “Henshin: advanced concepts and tools for in-place EMF model transformations,” *Model Driven Engineering Languages and Systems*, vol. 6394, pp. 121–135, 2010.
- [4] M. Lawley, K. Duddy, A. Gerber, and K. Raymond, “Language features for re-use and maintainability of mda transformations,” in *Workshop on Best Practices for Model-Driven Software Development*, 2004.
- [5] The Eclipse Foundation. (2014, May) Eclipse Modeling Framework project. [Online]. Available: <http://www.eclipse.org/modeling/emf/>
- [6] J. Miller and J. Mukerji, “MDA guide,” Object Management Group, Tech. Rep., 2003.
- [7] S. Kent, “Model driven engineering,” in *IFM*, 2002, pp. 286–298.
- [8] *Object Constraint Language*, Object Management Group Std., Rev. 2.0, Jun. 2005.
- [9] I. Kurtev, “Adaptability of model transformations,” Ph.D. dissertation, University of Twente, Enschede, 2005. [Online]. Available: <http://doc.utwente.nl/50761/>
- [10] T. Mens and P. van Gorp, “A taxonomy of model transformation,” *Electronic Notes in Theoretical Computer Science*, vol. 152, no. 0, pp. 125 – 142, 2006. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1571066106001435>
- [11] K. Czarnecki and S. Helsen, “Feature-based survey of model transformation approaches,” *IBM Systems Journal*, vol. 45, no. 3, pp. 621–645, 2006.

- [12] *Meta Object Facility (MOF) 2.0 Query/View/Transformation Specification*, Object Management Group Std., Rev. 1.1, Jan. 2011.
- [13] F. Jouault and I. Kurtev, "Transforming models with ATL," in *Satellite Events at the MoDELS 2005 Conference*. Springer, 2006, pp. 128–138.
- [14] E. Visser, "Program transformation with Stratego/XT," in *Domain-Specific Program Generation*, ser. Lecture Notes in Computer Science, C. Lengauer, D. Batory, C. Consel, and M. Odersky, Eds. Springer Berlin / Heidelberg, 2004, vol. 3016, pp. 315–349. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-25935-0_013
- [15] L. Baresi and R. Heckel, "Tutorial introduction to graph transformation: A software engineering perspective," in *Proceedings of the First International Conference on Graph Transformation*, ser. ICGT '02. London, UK: Springer-Verlag, 2002, pp. 402–429. [Online]. Available: <http://dl.acm.org/citation.cfm?id=647562.730670>
- [16] F. Budinsky, D. Steinberg, M. Paternostro, and E. Merks, *EMF: Eclipse Modeling Framework*, 2nd ed., ser. Eclipse series. Addison-Wesley, 2009.
- [17] A. Rensink, "The edge of graph transformation – graphs for behavioural specification," in *Graph Transformations and Model-Driven Engineering*, ser. Lecture Notes in Computer Science, G. Engels, C. Lewerentz, W. Schäfer, A. Schürr, and B. Westfechtel, Eds. Springer Berlin / Heidelberg, 2010, vol. 5765, pp. 6–32. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-17322-6_2
- [18] A. Corradini, U. Montanari, F. Rossi, H. Ehrig, R. Heckel, and M. Löwe, *Algebraic approaches to graph transformation. Part I: basic concepts and double pushout approach*. River Edge, NJ, USA: World Scientific Publishing Co., Inc., 1997, pp. 163–245. [Online]. Available: <http://dl.acm.org/citation.cfm?id=278918.278928>
- [19] T. Staijen, "Graph-based specification and verification for aspect-oriented languages," Ph.D. dissertation, University of Twente, Enschede, June 2010, iPA Dissertation Series ; 2010-04. [Online]. Available: <http://doc.utwente.nl/71550/>
- [20] A. Rensink, "The GROOVE simulator: A tool for state space generation," in *Applications of Graph Transformations with Industrial Relevance (AGTIVE)*, ser. Lecture Notes in Computer Science, J. Pfaltz, M. Nagl, and B. Böhlen, Eds., vol. 3062. Berlin: Springer Verlag, 2004, pp. 479–485. [Online]. Available: <http://doc.utwente.nl/66357/>
- [21] (2014, May) The Henshin project. Eclipse Foundation. [Online]. Available: <http://www.eclipse.org/modeling/emft/henshin/>
- [22] C. Ermel, E. Biermann, J. Schmidt, and A. Warning, "Visual modeling of controlled EMF model transformation using HENSHIN," *ECEASST*, vol. 32, 2010.
- [23] E. Syriani and H. Vangheluwe, "De-/re-constructing model transformation languages," *ECEASST*, vol. 29, pp. 1–14, 2010.

-
- [24] The Eclipse Foundation. (2014, May) Eclipse platform. [Online]. Available: <http://www.eclipse.org/>
 - [25] (2014, May) EMFText concrete syntax mapper. [Online]. Available: <http://www.emftext.org>
 - [26] D. Wagelaar. (2014, May) SimpleGT graph transformation language for EMFTVM. [Online]. Available: <http://code.google.com/a/eclipselabs.org/p/simplegt/>
 - [27] Eclipse Foundation. (2014, May) ATL modeling zoo. [Online]. Available: <http://www.eclipse.org/atl/atlTransformations/>