

Towards Automatic Model Synchronization from Model Transformations

Yingfei Xiong¹, Dongxi Liu¹, Zhenjiang Hu¹, Haiyan Zhao², Masato Takeichi¹ and Hong Mei²

¹Department of Mathematical Informatics
Graduate School of Information Science and Technology
University of Tokyo, Hongo 7-3-1, Bunkyo-ku, Tokyo 113-8656, Japan
{Yingfei_Xiong,liu,hu,takeichi}@mist.i.u-tokyo.ac.jp

²Institute of Software
School of Electronics Engineering and Computer Science
Peking University, Beijing, 100871, China
{zhhy,meih}@sei.pku.edu.cn

ABSTRACT

The metamodel techniques and model transformation techniques provide a standard way to represent and transform data, especially the software artifacts in software development. However, after a transformation is applied, the source model and the target model usually co-exist and evolve independently. How to propagate modifications across models in different formats still remains as an open problem.

In this paper we propose an automatic approach to synchronizing models that are related by model transformations. Given a unidirectional transformation between metamodels, we can automatically synchronize models in the metamodels by propagating modifications across the models. We have implemented our approach on the Atlas Transformation Language (ATL) and have tested our implementation on several ATL transformations.

1. INTRODUCTION

Model transformations play an important role in Model-driven architecture(MDA), an approach to software development, which provides a way to organize and manage software artifacts by automated tools and services for both defining models and facilitating transformations between different model types. Writing model transformations is becoming a common task in software development.

ATL [13] is a practical model transformation language that has been designed and implemented by INRIA to support specifying model transformations that can cover different domains of applications [1]. As a simple running example which will be used throughout this paper, consider the following UML2Java transformation in ATL:

```
module UML2Java;
```

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

```
create OUT : Java from IN : UML;

rule Class2Class {
  from u : UML!Class (
    not u.name.startsWith('__draft__')
  )
  to j : Java!Class (
    name <- u.name,
    fields <- u.attrs
  )
}

rule Attribute2Field {
  from a : UML!Attribute
  to f : Java!Field (
    name <- '_' + a.name,
    type <- a.type
  )
}
```

which uses two rules to transform a simple Unified Modeling Language (UML) model to a simple Java model. Roughly speaking, it maps each UML class whose name does not start with “_draft_” to a Java class with the same name, and each attribute of the class to a field of the corresponding Java class where the field name is the attribute name with an additional prefix “_”. For instance, this transformation maps the UML model (in XMI [21])

```
<Class name="Book" description="a demo class">
  <attrs name="title" type="String"/>
  <attrs name="price" type="Double"/>
</Class>
<Class name="__draft__Authors"/>
```

to the following Java model (in XMI).

```
<Class name="Book">
  <fields name="_title" type="String"/>
  <fields name="_price" type="Double"/>
</Class>
```

Despite a bunch of interesting applications of model transformations in software development, there is little work on a systematic method to maintain models at different stages of the software development. Models may be changed after the transformation in both the source side and the target side. For the above example, suppose a group of designers and a group of programmers are working on the models at the same time. The designers may add a new attribute **authors** to the **Book** class on the UML model

```
<Class name="Book" description="a demo class">
```

```

    <attrs name="title" type="String"/>
    <attrs name="price" type="Double"/>
    <attrs name="authors" type="String"/>
  </Class>
<Class name="__draft__Authors"/>

```

while at the same time the programmers may change the field name `_title` to `_bookTitle`, delete the field `_price` from the Java model, and add a new comment to the `Book` class.

```

<Class name="Book">
  <fields name="_bookTitle" type="String"/>
  <comment description="_price should be >= 0"/>
</Class>

```

Now the UML model and the Java model become inconsistent and need to be synchronized. Simply performing the transformation from UML model to Java model again is not adequate because the modifications on the Java model will be lost.

There are many challenges in automatically synchronizing these two models related by a model transformation. First and most importantly, in order to establish and maintain consistency, we need to precisely define what it means for two models to be synchronized. Although there are several general model synchronization frameworks [11, 12, 6] and many specific code-model synchronization tools such as Rational Rose [22], there is, as far as we are aware, no clear semantics for model synchronization in the context where a model transformation is formally given.

Second, we need an automatic way to derive from a given transformation enough necessary information, forward and backward, such that not only modifications on the source model can be automatically propagated to the target model, but also modifications on the target model can be automatically reflected back to the source model. The existing model synchronization systems [22] and model synchronization frameworks [11, 12, 6] cannot work well here, because they require users to explicitly write synchronization code to deal with each type of modification on each type of model. This makes it hard to guarantee consistency between the synchronization code and the transformation code, let alone to say consistency between the two models.

Third, our method is expected to be able to deal with general model transformations described in general transformation languages. That is, we do not expect our method to be only able to deal with a subset of transformations written in an expressive language or only able to deal with a transformations written in a restrictive language. In fact, the more restriction we impose on a model transformation, the easier but less useful the derived model synchronization process will be. Therefore, we should target a class of practically useful model transformations in order to obtain a useful model synchronization system.

In this paper, we report our first attempt towards automatically constructing a model synchronization system from a given model transformation described in ATL. Our main contributions can be summarized as follows.

- We define a clear semantics of model synchronization under the context where two models to be synchronized are related by a model transformation. Our semantics precisely characterizes the behavior of the synchronization process with four important properties, namely stability, information preservation, modification propagation and composability, which provide

users a clear image of what models will be after synchronization. These properties were much motivated by studies on updating semantics of database views [5] and the well-definedness of bidirectional tree transformation [9, 16]. We are the first who adapted these results to solve the model synchronization problem.

- We propose a new model synchronization approach that can automatically synchronize two models related by a transformation described in ATL, without requiring users to write extra synchronizing code. The model synchronization process satisfies the required properties and ensures the models will be correctly synchronized. Different from the existing bidirectional tree transformations working on high level functional programs [9, 16], our approach works on low level byte codes, which allows us to target more general transformation programs and cover the full ATL.
- We have implemented a model synchronization system by extending the ATL Virtual Machine (VM), the interpreter of ATL byte-code, and have successfully tested several ATL transformation examples in the ATL web site [1]. The current prototype system is available at our web site¹.

The rest of the paper is organized as the follows. We start by defining semantics of model transformation of two models that are related by a model transformation in Section 2. We then show how to automatically synchronize models related by a transformation in Section 3 and Section 4. We give a case study to illustrate the feasibility of our system in Section 5. Finally, we discuss related work in Section 6 and conclude the paper in Section 7.

2. SEMANTICS OF MODEL SYNCHRONIZATION

We consider synchronization of two models that are related by a model transformation. The semantics of model synchronization characterizes the behavior of the synchronization process. A well-defined semantics offers users clear information on what their models should be after synchronization. This will increase the confidence of users to deploy automatic model synchronization in practical software development.

As stated in Section 1, the semantics of model synchronization is described by a set of properties. Before describing the properties, we first give some notions and operators on model.

2.1 Model and Synchronization

A model is a function mapping from model references to model elements, and each model element is a function mapping from attribute names to values. For example, in an XMI file, model references are XLinks referring to some XML elements. The UML model in Section 1 maps `/0` and `/1` to two `Class` model elements, and these two model elements map the attribute name `name` to the value `"Book"` and `"__draft__authors"`, respectively.

The values in attributes are functions mapping from value addresses to single values. A single value can be a model reference, a `null` value or a value of boolean, string or integer.

¹<http://xiang.yingfei.googlepages.com/modelsynchronization>

A `null` value means an undefined value. Note in MetaObject Facility (MOF) Specification [19], attributes can store single values or different types of collections. For convenience of presentation, we abstract them all as functions mapping from addresses to single values because in actual implementation all values need to be accessed through memory addresses. Models can be constrained by a metamodel. In other words, a metamodel includes a set of models sharing the same structure constraints.

There are several notations to be used in the following presentation. The notation m is used for denoting a model, r for a model reference, n for an attribute name, d for a value address and v for a single value or a model element. Given two metamodels S and T , the model transformation $f : S \rightarrow T$ is a partial function that takes models in S and produces models in T .

In our approach, a *synchronization process* with respect to a given transformation $f : S \rightarrow T$ is a partial function with the following signature: $sync_f : S \times S \times T \rightarrow S \times T$ which takes as input the original source model, the modified source model and the modified target model, and produces the synchronized source model and target model. Note that $sync_f$ does not need the original target model since it can be obtained by transforming the original source model by f .

2.2 Modification Operators on Model

Figure 1 defines three operators `replace`, `delete` and `insert`, which perform replacement, deletion and insertion to models respectively, as indicated by their names. In this section, these operators help define the semantics of model synchronization; in the next section, they are used by the extended transformation system to implement the putting-back functions. These operators may take two or three parameters. The first parameter m is the model to be modified. The second parameter obj specifies the location in the model where a value or a model element is to be modified. The parameter obj can be r specifying a model element, can be (r, n) specifying an attribute in a model element, or can be (r, n, d) specifying a value in an attribute of a model element. The third parameter v , if available, is a new value for the element or attribute. In the definition of these operators, the notation $f[k \mapsto v]$ means a function that maps k to v and maps any other value $k' \neq k$ to $f(k')$. Also we use the notation \perp to indicate a function is undefined at the spot.

Suppose M is the collection of model elements. We define a modification operation ϕ as a function $\phi : M \rightarrow M$ implemented using one of the above operators. For instance, the modification operation ϕ defined by $\phi(m) = \text{delete}(m, r)$ denotes deleting the model element referred by r . Users may modify a model in many places at one time. This is modeled by a sequence ψ of modification operations, represented as $\phi_1 \circ \phi_2 \circ \dots \circ \phi_n$.

If two modification operations affect different parts of a model, they are said *distinct*, defined as below. The underline symbol $_$ below indicates the “don’t care” parameter.

DEFINITION 1. Let $op_1, op_2 \in \{\text{replace}, \text{delete}, \text{insert}\}$. Then two operations $op_1(m, obj_1, _)$ and $op_2(m, obj_2, _)$ are *distinct* if $obj_1 \neq obj_2$.

For distinct modification operations in a sequence, we can change their order without affecting modification result since they affect different parts of a model. On the other hand, the

sequence of non-distinct operations can be safely suppressed by taking only the last one. For example, if an operation to change the attribute of a model element into “a” is followed by another operation to change the same attribute into “b”, then the second operation can represent this sequence. Due to this, we consider only sequences with distinct modification operations in the following presentation, and assume Ψ denotes the set of all such sequences.

A modification operation is not distinct to itself. So if we apply the same sequence of operations twice, the operations at the first time can be suppressed completely. Thus ψ , a sequence of modification operations, is idempotent:

$$\forall \psi \in \Psi, \text{ and a model } m. \psi(\psi(m)) = \psi(m)$$

which can be used to check if a modification sequence ψ has been applied to a model or not. If we apply ψ to the model and the model remains the same, then ψ has already been applied to the model.

Some modifications to one model cannot be propagated to the other model. For the UML2Java example, the modification to the `comment` attribute in the Java model cannot be propagated to the UML model since there is no corresponding attribute. The operations performing such modifications are said to be *non-reflectable*, and otherwise they are *reflectable*.

DEFINITION 2. Given a transformation $f : S \rightarrow T$. A modification operation ϕ_t is *reflectable* w.r.t f if for any $s \in S$, there exists a modification operation ϕ_s such that $f(\phi_s(s)) = \phi_t(f(s))$.

2.3 Properties of Synchronization

The semantics of model synchronization is described by four properties in this section: *stability*, *preservation*, *propagation* and *composability*. One technical contribution of this paper is to adapt these properties from the areas of view update and bidirectional tree transformations[5, 9] to the area of model synchronization. In the following, we will describe what each of these properties means for model synchronization, and discuss the relations between them and the corresponding properties in the literature [5, 9]. Note that all these properties apply only when the execution of $sync_f$ process is successful.

The *stability* property says if neither of the source model and the target model are modified after transformation, the synchronization process should not modify any of them.

PROPERTY 1 (STABILITY). $sync_f(s, s, f(s)) = (s, f(s))$

The *stability* property corresponds to the `GETPUT` property [9] and the acceptable condition [5].

The *preservation* property states that the synchronization process should keep the modifications to source models and target models in the synchronized source models and the synchronized target models, respectively.

PROPERTY 2 (PRESERVATION). *Given $f : S \rightarrow T$, $s \in S$, $\psi_s, \psi_t \in \Psi$. If $sync_f(s, \psi_s(s), \psi_t(f(s))) = (s', t')$, then $\psi_s(s') = s'$ and $\psi_t(t') = t'$.*

By this property, for the UML2Java example in Section 1, programmers can expect their modifications on `comment` and `bookTitle` are kept on the Java model after the synchronization, while designers can expect the `authors` attribute still appears on the UML model. The *preservation* property

$$\begin{aligned}
\text{replace}(m, \text{obj}, v) &= \begin{cases} m[r \mapsto m(r)[n \mapsto m(r)(n)[d \mapsto v]]], & \text{if } \text{obj} = (r, n, d) \text{ and } m(r)(n)(d) \text{ is defined;} \\ m, & \text{otherwise.} \end{cases} \\
\text{delete}(m, \text{obj}) &= \begin{cases} m[r \mapsto m(r)[n \mapsto m(r)(n)[d \mapsto \perp]], & \text{if } \text{obj} = (r, n, d); \\ m[r \mapsto \perp], & \text{if } \text{obj} = r; \\ m, & \text{otherwise.} \end{cases} \\
\text{insert}(m, \text{obj}, v) &= \begin{cases} m[r \mapsto m(r)[n \mapsto m(r)(n)[d \mapsto v]]], & \text{if } \text{obj} = (r, n, d) \text{ and } m(r)(n)(d) \text{ is undefined;} \\ m[r \mapsto v], & \text{if } \text{obj} = r; \\ m, & \text{otherwise.} \end{cases}
\end{aligned}$$

Figure 1: Operators on Models

gets inspired by the PUTGET property [9] and the consistent condition [5], but these existing properties are defined in the situation where only views can be modified and thus concerns only preservation of modification to views.

The *propagation* property guarantees the correct propagation of modifications among models. That is, the synchronized target model t' contains all those modifications in ψ_s if they are applied to values used by transformation f , and the synchronized source model s' contains all reflectable modifications in ψ_t .

PROPERTY 3 (PROPAGATION). *Given $f : S \rightarrow T$, $s \in S$, $\psi_s, \psi_t \in \Psi$. If $\text{sync}_f(s, \psi_s(s), \psi_t(f(s))) = (s', t')$, then $\psi_t'(f(s')) = t'$, where ψ_t' consists of all non-reflectable modification operations in ψ_t .*

The rationale behind this property is that if one reflectable modification in ψ_t is not in s' , then it cannot be generated by applying f to s' , and thus the equation $\psi_t'(f(s')) = t'$ cannot hold; if one modification in ψ_s that will be brought into the target model by f is not in t' , then t' cannot equal $\psi_t'(f(s'))$ because $f(s')$ includes this modification. This property also gets inspired by the PUTGET property [9] and the consistent condition [5]. However, this property concerns two-way propagation of modifications and allows non-reflectable modifications on target models.

The last property we consider is *composability*. Intuitively, this property says synchronizing twice with two sequences of operations will have the same effect as synchronizing one with one sequence of operations that is concatenated from the two sequences of operations.

PROPERTY 4 (COMPOSABILITY). *Given $f : S \rightarrow T$, $s \in S$, $\psi_s, \psi_{s'}, \psi_t, \psi_{t'} \in \Psi$. If $\text{sync}_f(s, \psi_s(s), \psi_t(f(s))) = (s', t')$, $\text{sync}_f(s, \psi_{s'}(s'), \psi_{t'}(t')) = (s'', t'')$ and $\text{sync}_f(s, \psi_{s'}(\psi_s(s)), \psi_{t'}(\psi_t(f(s)))) = (s''', t''')$, then $(s'', t'') = (s''', t''')$.*

This property corresponds to the PUTPUT property [9] and gives users the freedom of performing synchronization at the time they want.

3. BACKWARD PROPAGATION OF MODIFICATIONS

To synchronize two models related by a model transformation, we need to propagate modifications between the source model and the target model. The propagation of modifications from the source model to the target model, i.e. the forward propagation, can be carried out by running the model transformation again. However, the propagation of modifications from the target model to the source model, i.e. the backward propagation, cannot get direct help from this transformation.

Table 1: The Core Instructions of ATL Byte-code

instructions	description
push	push a constant
pop	pop the top of the stack
store	store a value into a local variable
load	load value from local variable
if	branch if the top of the stack is <i>true</i>
iterate	delimitate the beginning of iteration on collection elements
enditerate	delimitate the end of iteration on collection elements
call	call a method
new	create a new model element
get	fetch an attribute of a model element
set	set an attribute of a model element

In this section, we will propose a technique to implement the backward propagation by extending the ATL Virtual Machine (VM). If we execute a transformation on this extended ATL VM, we will get a target model with extended model elements and extended single attribute values, and also a set of validity-checking functions. Extended model elements and extended single values contain putting-back functions, which help reflect back modifications to this value back into the corresponding value in the source model. The validity-checking functions are used to check, after backward propagation, whether the modified values in the source model are valid in terms that they do not change the execution path of the transformation over this updated source model. This is to guarantee that the preservation property is satisfied by our model synchronization process.

3.1 ATL Byte-code

An ATL transformation program is first compiled into ATL byte-code and then executed on the ATL VM. The ATL VM, like the Java virtual machine, contains a stack to hold local variables and partial results. An ATL byte-code program consists of a sequence of instructions. A summary of the core ATL instructions is given in Table 1. The full specification of ATL byte-code and the ATL virtual machine can be found at the ATL web site [1].

As a simple example, the rule `Attribute2Field` in the UML2Java transformation in the introduction can be written in byte-code, as shown in Figure 2. The first three lines return a list containing all `UML!Attribute` instances in the source model. Then instructions between Line 4 and Line 19 iterate on the list. Each instance is stored in a variable `a` (Line 5) and for each instance, a `Java!Field` model element is created (Line 6-7) and stored in a variable `f` (Line 8). Then the `name` attribute of the variable `a` is concate-

```

1  push "UML!Attribute"
2  push "IN"
3  call "S.allInstancesFrom(S):QJ"
4  iterate
5  store "a"
6  push "Java!Field"
7  new
8  store "f"
9  load "f"
10 push "_"
11 load "a"
12 get "name"
13 call "S.Concatenate(S):S"
14 set "name"
15 load "f"
16 load "a"
17 get "type"
18 set "type"
19 enditerate

```

Figure 2: Byte-code for Attribute2Field

nated with “_” (Line 10-13) and set to the **name** attribute of the variable **f** (Line 9 and 14). The **type** attribute of the variable **a** is retrieved (Line 16 and 17) and set to the **type** attribute of the variable **f** (Line 15 and 18).

3.2 Extending the ATL Virtual Machine (VM)

In the extended VM, each model element and each single value are associated with a set of putting-back functions: **rep**, **del**, **sat_r**, **sat_d** and **val**. The function **rep** is to be called when the single value is replaced, the function **del** is to be called when the single value or the model element is deleted, the function **sat_r** is used to check whether the replacement is valid to be put back, the function **sat_d** is used to check whether the deletion is valid to be put back, and the function **val** is used to reevaluate the single value or the model element from the source model.

Specifically, we made three extensions to the ATL VM. The first is that the model elements or single values in source models are extended with putting-back functions when the models are loaded. The second extension is to extend the semantics of each ATL byte-code instruction, which, if generating new values, also associates the generated values with appropriate putting-back functions. In addition, each **if** instruction also generates a validity-checking function to ensure that its condition is still satisfied after propagating modifications into source models. The third extension is made on the ATL library methods, such as **Concatenate** and **startsWith**, such that the values returned by these methods are also associated with putting-back functions. In most methods and some instructions, new values are created by composing existing values. In those cases, the putting-back functions of new values are built by composing the putting-back functions of existing values. In this way, a call to a putting-back function of a new value will invoke a series of calls to functions of existing values, and will eventually calls putting-back functions of source model values to update the source model if necessary. herefore when a model element or a single value in the target model is modified (replaced or deleted), we can call appropriate putting-back functions to propagate the modification back into the source model.

3.2.1 Extending Source Models

The model elements and single attribute values in source models are extended before transformations. This is done when the ATL VM loads source models into its runtime

environment.

Suppose v is a single value at the location of $m(r)(n)(d)$. Then its extension is represented as (v, ext) , where $ext = (\mathbf{rep}, \mathbf{del}, \mathbf{sat_r}, \mathbf{sat_d}, \mathbf{eval})$ and each function in this tuple is defined as below with the operators in Figure 1.

$$\begin{aligned}
 \mathbf{rep}(m, v') &= \mathbf{replace}(m, (r, n, i), v') \\
 \mathbf{del}(m) &= \mathbf{delete}(m, (r, n, i)) \\
 \mathbf{sat_r}(v') &= \mathbf{true} \\
 \mathbf{sat_d}() &= \mathbf{true} \\
 \mathbf{val}(m) &= m(r)(n)(i)
 \end{aligned}$$

Here the function **rep** and **del** replace and delete the value in the source model, respectively. The **sat_r** and **sat_d** functions are always true, meaning that the associated value can always be replaced or removed. The **val** function just returns the value from the source model.

The extension to the model element $v = m(r)$ is represented as (v, ext) , where $ext = (\mathbf{rep}, \mathbf{del}, \mathbf{sat_r}, \mathbf{sat_d}, \mathbf{val})$. These functions are defined as below.

$$\begin{aligned}
 \mathbf{rep}(m, v') &= m \\
 \mathbf{del}(m) &= \mathbf{delete}(m, r) \\
 \mathbf{sat_r}(v') &= \mathbf{false} \\
 \mathbf{sat_d}() &= \mathbf{true} \\
 \mathbf{val}(m) &= m(r)
 \end{aligned}$$

These functions have the same meaning as the above ones. Note that a model element cannot be replaced, so the **rep** function does nothing and the **sat_r** always returns false.

3.2.2 Extending ATL Byte-code Instructions

Some instructions of ATL byte-code do not change or create values or model elements, but move values among different parts (e.g. from models to the stack) of the running environment. The instructions **pop**, **store**, **load**, **get** in Table 1 belong to this case. We extend these instructions so that they not only move the original value but also the putting-back functions. Although the **set** instruction modifies model elements, we treat it as an instruction moving a value from the stack to a model and extend the **set** instruction in the same way.

In the following, we explain how to extend the instructions **push**, **iterate**, **enditerate**, **new** and **if**. The **call** instruction is discussed in the next subsection.

push *cst*

The original semantics of this instruction is to push the constant *cst* onto the top of the operand stack. For example, the instruction at line 10 in Figure 2 pushes a constant string ‘_’ to the stack. In the extended ATM VM, the system pushes a extended constant (cst, ext) , where $ext = (\mathbf{rep}, \mathbf{del}, \mathbf{sat_r}, \mathbf{sat_d}, \mathbf{val})$, and these putting-back functions are defined as below.

$$\begin{aligned}
 \mathbf{rep}(m, v') &= m \\
 \mathbf{del}(m) &= m \\
 \mathbf{sat_r}(v') &= \mathbf{if } v' = cst \mathbf{ then } true \mathbf{ else } false \\
 \mathbf{sat_d}() &= \mathbf{false} \\
 \mathbf{val}(m) &= cst
 \end{aligned}$$

Since the modifications on *cst* cannot be reflected back to the source model, we do not allow the replacement or deletion of this value. So the **rep** and **del** functions do nothing; the **sat_r** and **sat_d** functions are always false.

new, **iterate** and **enditerate**

The `new` instruction creates new target model elements. However, this instruction provides no information of what source model element or source value corresponds to the new target model element.

To create a collection of target model elements, usually we have to traverse a collection of values or model elements, and create a target model element for each item in the collection. Thus items in the collection can be consider as sources of the target model elements. For the example in Figure 2, a set of `Field` model element is created when iterating over the set of `Attribute` model elements in the source. In ATL byte-code the only way to traverse a collection is the `iterate` and `enditerate` instructions.

Based on the above observation, we create a stack `IterObjs` in the runtime environment to remember the objects being iterated. The `iterate` instruction pushes the object being iterated onto the `IterObjs` stack, while the `enditerate` instruction pops off the top object from the `IterObjs` stack. For the model element created by the `new` instruction, it copies the putting-back functions from the object at the top of the `IterObjs` stack. That is, suppose the object at the top of the `IterObjs` stack is (o, ext) . Then a model element v created by the instruction `new` will be associated with the extension ext and becomes (v, ext) . If a model element is created outside any iteration, it is considered as a constant and the putting-back functions for constants are associated to the model element.

`if l`

The `if` instruction jumps to the instruction with label l if the value at the top of the operand stack is true, otherwise it falls through to the next instruction. We call the value of at the top of the stack the *condition value* of the `if` instruction. If we execute the transformation again after backward propagation of modifications, some condition values may become different from their values before backward propagation. This will change the execution paths of the transformation, and probably generate target models in which the user modifications are lost. In our synchronization framework, this will violates the preservation property.

In our running example, a `Java!Class` model element is generated only when the `name` attribute of the `UML!Class` model element does not start with `__draft__`. Suppose a user happens to change the `name` attribute of a `Java!Class` model element to a value starting with `__draft__`. After propagating modifications backward and transforming again, this model element will disappear on the target model.

To prevent such cases, we require that modifications by users should not cause a condition value to be different before and after backward propagation. Our solution is that when executing an `if` instruction, the system will generate a validity-checking function `sat_c`, and store the function into a set Θ . After backward propagation, this validity-checking function is used to recompute the condition value of this `if` instruction and check if it is the same as before backward propagation. If not, the system reports an error.

Suppose when executing an `if` instruction, its condition value is (v, ext) , where $ext = (\text{rep}, \text{del}, \text{sat}_r, \text{sat}_d, \text{val})$. Then the function `sat_c` generated for this `if` instruction is: $\text{sat}_c(m) = \text{if } \text{val}(m) = v \text{ then true else false}$.

After backward propagation, the system calls all validity-checking functions in Θ and report a failure if a function returns false.

3.2.3 Extending ATL Library Methods

The `call` instruction is to call ATL library methods. These methods are implemented in Java, not ATL byte-code, so we need to extend them to return extended model elements or extended single values. In the following, we will explain how to extend ATL library methods `Concatenate` and `startsWith` as examples.

The methods `Concatenate` and `startsWith` both take as arguments the first two strings at the top of the operand stack. Suppose the two arguments for both `Concatenate` and `startsWith` methods are (str_1, ext_1) and (str_2, ext_2) , where $ext_1 = (\text{rep}_1, \text{del}_1, \text{sat}_r1, \text{sat}_d1, \text{val}_1)$ and $ext_2 = (\text{rep}_2, \text{del}_2, \text{sat}_r2, \text{sat}_d2, \text{val}_2)$.

For the concatenated string returned by the `Concatenate` method, its putting-back functions are $(\text{rep}, \text{del}, \text{sat}_r, \text{sat}_d, \text{val})$, as defined below:

```

rep(m, v') = repx(m, v', 0)
repx(m, v', i) =
  if sat_r1(head(v', i)) and sat_r2(tail(v', len(v') - i)) then
    rep1(rep2(m, tail(v', len(v') - i), head(v', i)))
  else repx(m, v', i + 1)
del(m) = if sat_d1(m) then del1(delx(m))
delx(m) = if sat_d2(m) then del2(m)
sat_r(v') = sat_rx(v', 0)
sat_rx(v', i) =
  if i ≤ len(v') then
    if sat_r1(head(v', i)) and sat_r2(tail(v', len(v') - i))
    then true
    else sat_rx(v', i + 1)
  else false
sat_d() = sat_d1() or sat_d2()
val(m) = val1(m) ⊕ val2(m)

```

The function `tail(v', l)` extracts the tail substring of string v' of length l ; the function `head(v', l)` extracts the leading substring of string v' of length l . The operator \oplus is used to concatenate two strings in the above definition. These putting-back functions ensure a reasonable putting-back behavior so long as strings are separated with constants.

For the boolean value returned by the `startsWith` method, its putting-back functions are defined as below.

```

rep(m, v') = m
del(m) = m
sat_r(v') = false
sat_d() = false
val(m) = substr(val1(m), val2(m))

```

Boolean values returned by the `startsWith` method cannot be modified, but these values can be reevaluated by calling the `val` function. The `substr` checks whether the first argument is the leading substring of the second argument.

4. SYNCHRONIZATION FRAMEWORK

In this section we show how to realize our model synchronization process (as defined in Section 2)

$$\text{sync}_f : S \times S \times T \rightarrow S \times T$$

based on (1) a given transformation $f : S \rightarrow T$ which shows how to map the source model (including its modification) to the target model, and (2) the derived putting-back functions (in Section 3) which shows how to reflect modifications (replacements and deletions) on target models back to source models. We shall illustrate our synchronization algorithm

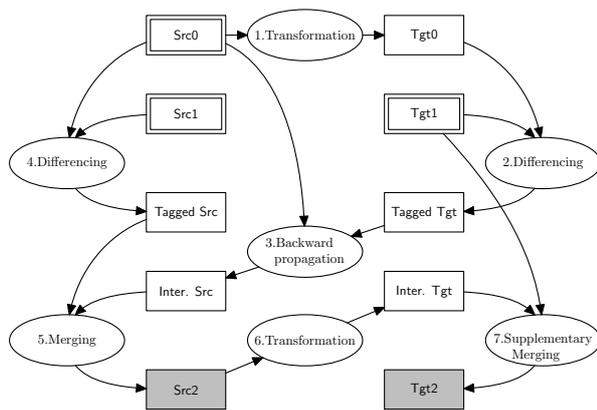


Figure 3: Overview of Synchronization Algorithm

by our running example, and explain intuitively that our synchronization satisfies the properties in Section 2.

4.1 Synchronization Algorithm

An overview of our synchronization algorithm is shown in Figure 3. The synchronization algorithm takes as input

- the original source model `Src0`,
- the modified source model `Src1`,
- the modified target model `Tgt1`, and
- the transformation f which can generate a target model from a source model

and returns as output

- the synchronized source model `Src2`, and
- the synchronized target model `Tgt2`.

It should be noted that our synchronization algorithm makes use of the original source model. This is in sharp contrast to other systems [3, 7, 4], and it contributes much to the good properties of our system (see Section 4.2).

The basic idea of the algorithm is: first put back the modifications on the target into the source and merge with modifications on the source, then reproduce the target model. The synchronization process in all has seven steps, which will be informally illustrated through our running example of UML2Java in Section 1, where all the inputs have been given.

Step 1: Generating the original target model

This step simply applies the transformation to the original source model to obtain the original target model `Tgt0`. For our UML2Java example, it is the first Java model in the introduction.

Step 2: Deriving modified target model with modification tags

We use modification tags to indicate the modifications users have performed on models. Modification tags can be annotated on primitive values and on model elements, and are defined below:

$\text{ModTag} = \{\text{Non}, \text{Rep}, \text{Ins}, \text{Del}\}$

The tag `Non`, often being omitted, indicates a value or a model element has not been modified. The tag `Rep` indicates a primitive value has been replaced by another primitive value. The tag `Ins` indicates a model element or a primitive value in a collection is inserted by users. The tag

`Del` indicates a model element or a primitive value in a collection is deleted by users.

We apply the existing *differencing* algorithms [2, 17] to find what modifications that users have made on the target model. The differencing procedure compares the original model and the modified model, and produces a new model annotated with modification tags. Return to our running example, differencing the original Java model with the modified Java model yields the following tagged model `Tagged Tgt`, where modification tags are annotated as superscripts.

```

<Class name="Book">
  <fields nameRep="_bookTitle" type="String"/>
  <fieldsDel name="_price" type="Double"/>Del
  <commentRep description="_price should be >= 0"/>Ins
</Class>
  
```

It should be noted that adding the comment to the class needs two modifications. First a new `Comment` model element need to be inserted. We put this tag at the end of the model element. Second the `comment` attribute of the class need to be modified from `null` to the reference to the comment. We put this tag on the attribute name. The same tagging method are used for deleting the `_price` field.

Step 3: Reflecting modification on the target model back to the source model

We apply the technique described in Section 3 to put back all reflectable modifications annotated in the model `Tagged Tgt1` back to the source model, resulting in an updated model `Inter.Src` (i.e., an intermediate source model).

It is possible that multiple modifications are reflected to one value or one model element. In this case, the framework uses rules in Tables 2 and 3 to merge the modifications by comparing the modifications tags to be applied by the two modifications.

```

<Class name="Book" description="a demo class">
  <attrs nameRep="bookTitle" type="String"/>
  <attrsDel name="price" type="Double"/>Del
</Class>
<Class name="__draft__Authors"/>
  
```

Note that the inserted `comment` on the target model is not reflected to the source model. This is because the given transformation only relates UML classes with Java classes and attributes with fields.

Step 4: Deriving modified source model with modification tags

This step is similar to Step 2 except it is applied to the source model instead of the target model. Differencing the original source model `Src0` with the modified source model `Src1`, this step produces a tagged model `Tagged Src`.

```

<Class name="Book" description="a demo class">
  <attrs name="bookTitle" type="String"/>
  <attrs name="price" type="Double"/>
  <attrsIns name="authors" type="String"/>Ins
</Class>
<Class name="__draft__Authors"/>
  
```

Step 5: Merging two modified source models

Now `Tagged Src` contains modifications on the source model and `Inter.Src` contains modifications on the target model. Then the framework uses a *merging process* to merge the two models into one by comparing the modification tags according to the rules in Tables 2 and 3. After merging, the merged model `Src2` should have the modifications from both sides if there is no conflict. Otherwise, a conflict error should be reported.

Table 2: Rules for merging tagged values v_1 and v_2

$v_1.tag$	$v_2.tag$	condition	result
Non	-	-	v_2
Del	Del/Non	-	v_1
Del	Rep/Ins	-	conflict
Rep/Ins	Rep/Ins	$v_1 = v_2$	v_1
Rep/Ins	Rep/Ins	$v_1 \neq v_2$	conflict
Rep/Ins	Del	-	conflict
Rep/Ins	Non	-	v_1

Table 3: Rules for merging tagged model elements e_1 and e_2

$e_1.tag$	$e_2.tag$	result tag
Non	Del/Non/Ins	$e_2.tag$
Del	Del/Non	Del
Del	Ins	conflict
Ins	Ins/Non	Ins
Ins	Del	conflict

```

<Class name="Book" description="a demo class">
  <attrs nameRep="bookTitle" type="String"/>
  <attrsDel name="price" type="Double"/>Del
  <attrsIns name="authors" type="String"/>Ins
</Class>
<Class name="__draft__Authors"/>

```

Step 6: Propagating all modifications on the source model to the target model

In order to propagate the merged modifications to the target side, we apply the transformation on `Src2` and get `Inter.Tgt` (i.e. an intermediate target model).

```

<Class name="Book">
  <fields nameRep="_bookTitle" type="String"/>
  <fieldsDel name="_price" type="Double"/>Del
  <fieldsIns name="_authors"/>Ins
</Class>

```

Step 7: Supplementary merging on target models

`Inter.Tgt` now should contain the modifications on the source model and the reflectable modifications that has been reflected from the target model to the source. Yet the non-reflectable modifications are still missing. To merge such modifications, we copy the non-reflectable modifications from `Tgt1` and produce the synchronized target model `Tgt2`.

```

<Class name="Book">
  <fields nameRep="_bookTitle" type="String"/>
  <fieldsDel name="_price" type="Double"/>Del
  <fieldsIns name="_authors"/>Ins
  <commentIns description="\_price should be >= 0"/>Ins
</Class>

```

It is worth noting that to merge the modifications, we should first identify what modifications are non-reflectable and need to be merged. Here we define three types of identifiable non-reflectable modifications, as shown below:

- Replacements on attributes that have not been set during the transformation, e.g. the `comment` attribute on the `Java!Class` model element.
- Adding model elements of a type whose instance has never been created during the transformation, e.g., the new `Comment` element user added on the `Java` model.

- Adding references that refer to the model elements identified in the second type. For example, suppose there is another transformation that generates skeleton Java code from UML class diagrams. If later programmers add statements to Java methods, the references from the Java methods to the statements are non-reflectable and need to be copied.

All the three types of modifications can be identified by keeping track of what attributes have been set and what types of model elements have been created during the transformation.

4.2 Properties

It is worth remarking that our synchronization system satisfies the properties we described in Section 2.3. To prove it formally we need to give formal semantics to all ATL statements, and this formalization, however, is beyond the scope of the paper. So we only give an intuitive discussion of the properties.

First, our synchronization system satisfies the stability property. If users have not made modifications on models after transformation, our system will not put any modification tags on models so no models will be changed during synchronization.

Second, our synchronization system satisfies the preservation property. On the source side, the merge process will merge all modifications into the synchronized model `Src2`. On the target side, all reflectable modifications will be put back to the source. Because all condition expressions will be evaluated to the same value, all the reflectable modifications will be produced again following the same path. On the other hand, all non-reflectable modifications will be merged into the synchronized target model during the supplementary merging process. So modifications on the target model are preserved.

Third, our synchronization system satisfies the propagation property. This is directly followed from the last two steps of our synchronization process.

Finally, our synchronization system satisfies the composability property. Because the system ensures all condition expressions remains the same during synchronization, modifications are propagated in the same way regardless of how many times we synchronize.

5. A CASE STUDY

Our system has been successfully applied to several ATL examples listed at ATL web site [1]. In this section, we will use one of these examples to help demonstrate our approach described before. This example is about a transformation from class models to relational database models and is widely used in the literature of model transformations [15]. By this case study, we can see after users write an ATL transformation, the consistency of the source model and target model can be automatically maintained by our system when they are evolved, and the synchronization procedure exhibits some interesting properties.

To run this example, we need the ATL code, the source model as well as the source and target metamodels. Due to space limitation, only the source model is shown in Figure 4, and other files can be found at ATL web site [1]. This source model includes two classes `Person` and `Family`, and two Datatypes `String` and `Integer`. Each class has a

```

0: <?xml version="1.0" encoding="ISO-8859-1"?>
1: <xmi:XMI xmi:version="2.0" xmlns="Class"
   xmlns:xmi="http://www.omg.org/XMI" >
2:   <Class name="Person" ID="1">
3:     <attr name="firstName" ID="5" type="3"/>
4:     <attr name="closestFriend" ID="6" type="1"/>
5:     <attr name="emailAddresses" ID="7"
6:       multiValued="true" type="3"/>
7:   </Class>
8:   <Class name="Family" ID="2">
9:     <attr name="name" ID="8" type="3"/>
10:    <attr name="members" ID="9"
11:      multiValued="true" type="1"/>
12:   </Class>
13:   <DataType name="String" ID="3"/>
14:   <DataType name="Integer" ID="4"/>
15: </xmi:XMI>

```

Figure 4: A Source Model in XMI

collection of attributes `attr`, which can be single-valued or multi-valued. The attribute ID in each model element is added by us to identify model elements.

In this example, a class will be transformed into a table, and a datatype into a type in the relational table model. Each attribute in a class, if it is single-valued, will lead to a column in the corresponding table, otherwise a new table will be generated for it. And each table generated from a class also includes a key column. The ATL web site has the detailed description for this transformation. The target model generated by this transformation is given in Figure 5.

In the following, we will give several experiments to show the synchronization results of our system. Each experiment is to demonstrate some properties that our approach has.

In the first experiment, we invoke the synchronization procedure without changing the source model and the target model. After synchronization, the resulting source model and target model are still the same as the original ones, embodying the property of stability.

In the second experiment, change `Person.emailAddresses` in Line 14 to `Individual.emailAddresses` and change the type of `emailAddresses` in Line 17 from "3" to "4", that is, the type changes to `Integer`. In addition, we change the source model by removing the line 4, that is, the attribute of `closestFriendId` in class `Person` is deleted. After synchronization, the result source model keeps the attribute of `closestFriendId` deleted while the class name in line 2 changes from `Person` to `Individual` and the type of `emailAddresses` changes to "4", that is, changes to type `Integer`; the result target model has `closestFriend` originally in Line 10 deleted, the type of `emailAddresses` remaining `Integer` and all occurrences of the string "Person" changing to "Individual", in other words, the table name in Line 7 changes to `Individual`, the table name in Line 14 remains `Individual.emailAddresses`, and the column name in Line 15 changes to `IndividualID`. This experiment demonstrates the preservation property and propagation property.

In the third experiment, we change the string `objectId` in the line 8 into `objId`. This string comes from the transformation code, not from the source model. The system reports a failure during synchronization. This shows that our system has the ability to detect and report inappropriate modifications.

The fourth experiment is to demonstrate the composability property by dividing the modifications in the second experiment in two steps, that is, first change the table name

```

0: <xmi:XMI xmi:version="2.0" xmlns="Relational"
1:   xmlns:xmi="http://www.omg.org/XMI" >
2:   <Table name="Family" ID="2" key="1002">
3:     <col name="objectId" ID="1002" keyOf="2"
4:       type="4"/>
5:     <col name="name" ID="8" type="3"/>
6:   </Table>
7:   <Table name="Person" ID="1" key="1001">
8:     <col name="objectId" ID="1001" keyOf="1"
9:       type="4"/>
9:     <col name="firstName" ID="5" type="3"/>
10:    <col name="closestFriendId" ID="6"
11:      type="4"/>
11:   </Table>
12:   <Type name="String" ID="3"/>
13:   <Type name="Integer" ID="4"/>
14:   <Table name="Person_emailAddresses" ID="7">
15:     <col name="PersonId" ID="1007" type="4"/>
16:     <col name="emailAddresses" ID="1008"
17:       type="3"/>
18:   </Table>
19:   <Table name="Family_members" ID="9">
20:     <col name="FamilyId" ID="1009" type="4"/>
21:     <col name="membersId" ID="1010" type="4"/>
22:   </Table>
23: </xmi:XMI>

```

Figure 5: A Target Model in XMI

in the target model and delete the attribute in the source model, synchronize, then change the type of `emailAddresses` in the target model and synchronize again. After the two synchronization processes, we get the same result as the second experiment.

6. RELATED WORK

There have been a large number of approaches to model transformations, each with its own characteristics. To classify existing transformation approaches, Czarnecki et al. [8] have proposed a classification framework. This framework uses a set of features to classify model transformation approaches. Among them, bidirectionality is of great interest. This feature can be achieved through bidirectional languages that can be executed both forwardly and backwardly. The forward transformation takes the source model and produces the target model while the backward transformation takes the target model and produces the source model. Akehurst and Kent [3] use symmetric relationships to relate the source model and the target model symmetrically. Later this idea appears in some submissions [7, 4] to Query/View/-Transformation (QVT) Request for Proposal (RFP) and is adopted in the QVT final adopted specification [20]. Other researchers achieve bidirectionality by using existing graph transformation techniques like Triple Graph Grammars (TGGs) [10, 14]. Bidirectional languages, however, are not adequate to support synchronization because the transformations overwrite existing models without using information in the models. When a model is reproduced, information not presented in the other model will be lost. Furthermore, the source and the target models cannot be modified at the same time.

Our approach are much inspired by the studies of the view update problem in Database systems [5] and bidirectional tree transformations [16][9]. Bidirectional tree transformations, different from bidirectional model transformation languages, concern that after a transformation, how to propagate modifications on the target tree back to a existing source tree. We adapt their ideas to models, but different from these approaches which work on high-level functional programs, our approach work on low-level byte-code pro-

gram, which is promising to support more general transformations if they are written in languages that can be translated into ATL byte-code.

Most existing systems [22] and approaches [18] to model synchronization are specific, that is, they can only synchronize models in specific metamodels. Some researchers, however, have considered general model synchronization. Ivkovic and Kontogiannis [11] propose a general framework where modifications are represented as model transformations and synchronization is converting transformations on one model to transformations on the other model. Johann and Egyed [12] concern more about how to integrate synchronization into modeling tools, and propose a framework to incrementally synchronize models between model tools. Bottoni [6] uses graph rewriting rules to synchronize models. However, all the approaches require users to manually write rules to convert modifications on one model into modifications on the other model. In most cases, users have to write for each type of modification on each model, which is a considerable task. Compared to them, our approach extracts such information automatically from existing model transformations, and do not require users to write extra code.

The term “synchronization” sometimes refers to approaches to differencing and merging models in the same metamodel [2, 17]. These approaches can be used in our synchronization algorithm to difference and merge models.

7. CONCLUSION

In this paper we have reported our first attempt towards automatic construction of model synchronization systems under the condition that the models to be synchronized are related by model transformations. In our framework, if a model transformation from one model to another is given, these two models can be synchronized for free without writing extra code. The key contributions of our approach are two folds: an automatic derivation of putback codes from execution of a model transformation, and a new synchronization framework with clear synchronization semantics. We have implemented all the ideas in this paper as a system for synchronizing models transformed by ATL transformations. The experimental results are encouraging; several nontrivial examples in the ATL Web site have been successfully tested.

One limitation of our current system is that it cannot deal well with insertions on the target side; although the system works well on non-reflectable insertions on the target side, it cannot deal with reflectable insertions. We are solving this problem by introducing virtual holes to the source side. This is one of our future work.

8. REFERENCES

- [1] The ATL web site. <http://www.eclipse.org/m2m/atl/>.
- [2] M. Abi-Antoun, J. Aldrich, N. Nahas, B. Schmerl, and D. Garlan. Differencing and merging of architectural views. In *ASE '06: Proceedings of the 21st IEEE International Conference on Automated Software Engineering*, pages 47–58. IEEE Computer Society, 2006.
- [3] D. H. Akehurst and S. Kent. A relational approach to defining transformations in a metamodel. In *UML '02: Proceedings of the 5th International Conference on The Unified Modeling Language*, pages 243–258. Springer-Verlag, 2002.
- [4] B. K. Appukuttan, T. Clark, A. Evans, G. Maskeri, S. Reddy, P. Sammut, L. Tratt, R. Venkatesh, and J. S. Willans. QVT-Partners revised submission to QVT RFP, 2003.
- [5] F. Bancilhon and N. Spyrtos. Update semantics of relational views. *ACM Trans. Database Syst.*, 6(4):557–575, 1981.
- [6] P. Bottoni, F. Parisi-Presicce, S. Pulcini, and G. Taentzer. Maintaining coherence between models with distributed rules: from theory to Eclipse. In *GT-VMT '06: Proceedings of International Workshop on Graph Transformation and Visual Modeling Techniques*. Elsevier Science, 2006.
- [7] Compuware Corporation and SUN Microsystems. XMOF queries, views and transformations on models using MOF, OCL and patterns. <http://www.omg.org/docs/ad/03-08-07.pdf>, 2003.
- [8] K. Czarnecki and S. Helsen. Classification of model transformation approaches. In *OOPSLA 03 Workshop on Generative Techniques in the Context of Model-Driven Architecture*, 2003.
- [9] J. N. Foster, M. B. Greenwald, J. T. Moore, B. C. Pierce, and A. Schmitt. Combinators for bi-directional tree transformations: a linguistic approach to the view update problem. In *POPL '05: ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 233–246, 2005.
- [10] H. Giese and R. Wagner. Incremental model synchronization with triple graph grammars. In O. Nierstrasz, J. Whittle, D. Harel, and G. Reggio, editors, *Models '06: Proc. of the 9th International Conference on Model Driven Engineering Languages and Systems*, volume 4199, pages 543–557, 2006.
- [11] I. Ivkovic and K. Kontogiannis. Tracing evolution changes of software artifacts through model synchronization. In *ICSM '04: Proceedings of the 20th IEEE International Conference on Software Maintenance*, pages 252–261. IEEE Computer Society, 2004.
- [12] S. Johann and A. Egyed. Instant and incremental transformation of models. In *ASE '04: Proceedings of the 19th IEEE international conference on Automated software engineering*, pages 362–365. IEEE Computer Society, 2004.
- [13] F. Jouault and I. Kurtev. Transforming models with ATL. In *Proceedings of Satellite Events at the MoDELS 2005 Conference*, volume 3844 of *Lecture Notes in Computer Science*, pages 128–138. Springer, 2006.
- [14] A. Konigs and A. Schurr. Tool integration with triple graph grammars - a survey. *Electronic Notes in Theoretical Computer Science*.
- [15] M. Lawley, K. Duddy, A. Gerber, and K. Raymond. Language features for re-use and maintainability of MDA transformations. In *OOPSLA Workshop on Best Practices for Model-Driven Software Development*, 2004.
- [16] D. Liu, Z. Hu, and M. Takeichi. Bidirectional interpretation of xquery. In *PEPM '07: Proceedings of the 2007 ACM SIGPLAN symposium on Partial evaluation and semantics-based program manipulation*, pages 21–30. ACM Press, 2007.
- [17] A. Mehra, J. Grundy, and J. Hosking. A generic approach to supporting diagram differencing and merging for collaborative design. In *ASE '05: Proceedings of the 20th IEEE/ACM international Conference on Automated software engineering*, pages 204–213. ACM Press, 2005.
- [18] U. Nickel, J. Niere, J. Wadsack, and A. Zündorf. Roundtrip engineering with FUJABA. In *WSR '00: Proceedings of 2nd Workshop on Software-Reengineering*, 2000.
- [19] Object Management Group. Metaobject facility (mof) specification. www.omg.org/docs/formal/02-04-03.pdf, 2002.
- [20] Object Management Group. MOF QVT final adopted specification. <http://www.omg.org/docs/ptc/05-11-01.pdf>, 2005.
- [21] Object Management Group. XML metadata interchange (XMI) specification, v2.1. www.omg.org/docs/formal/05-09-01.pdf, 2005.
- [22] T. Quatrani. *Visual Modeling with Rational Rose 2002 and UML*. Addison-Wesley Longman Publishing Co., Inc., 2002.