

GRoundTram: An Integrated Framework for Developing Well-Behaved Bidirectional Model Transformations^{*2}

Soichiro Hidaka

National Institute of Informatics
hidaka@nii.ac.jp

Zhenjiang Hu

National Institute of Informatics
hu@nii.ac.jp

Kazuhiro Inaba

National Institute of Informatics^{*1}
kinaba@nii.ac.jp

Hiroyuki Kato

National Institute of Informatics
kato@nii.ac.jp

Keisuke Nakano

University of Electro-Communications
ksk@cs.uec.ac.jp

ABSTRACT

Bidirectional model transformation is useful for maintaining consistency between two models, and has many potential applications in software development including model synchronization, round-trip engineering, and software evolution. Despite these attractive uses, the lack of a practical tool supporting for systematic development prevents it from being widely used. In this paper, we solve this problem by proposing an integrated framework called GRoundTram (Graph Roundtrip Transformation for Models), which is carefully designed and implemented for compositional development of well-behaved and efficient bidirectional model transformations. GRoundTram is built upon a well-founded bidirectional framework and is equipped with a user-friendly language for coding bidirectional model transformations, a novel tool for validating both models and transformations, an optimization mechanism for improving efficiency, and a powerful debugging environment for testing bidirectional behavior. GRoundTram has been used by other research groups besides ourselves and their results show its usefulness in practice.

KEYWORDS

Model-driven Development, Bidirectional Transformation, Model Transformation, Graph Transformation

1 Introduction

Bidirectional model transformation [Antkiewicz and Czarnecki, 2007, Czarnecki et al., 2009, Ehrig et al., 2007, Stevens, 2008, 2010] is an enhancement of model transformation with bidirectional capability, and is an important requirement in Object Management Group (OMG)'s Queries/Views/Transformations (QVT) standard for defining model transformation languages. It describes not only a forward transformation from a source model to a target model, but also a backward transformation showing how to reflect the changes in the target model in the source model so that consistency between them is maintained. Bidirectional model trans-

formation has many potential applications in software development, including model synchronization [Antkiewicz and Czarnecki, 2007, Giese and Wagner, 2006, Xiong et al., 2007], round-trip engineering [Antkiewicz and Czarnecki, 2006], software evolution [Lämmel, 2004], and multiple-view software development [Garcia, 2008, Grundy et al., 1998].

Unlike (unidirectional) model transformation where lots of tools have been developed for supporting design, validation, and test of model transformation, bidirectional model transformation lacks such useful tools, and that prevents it from being widely used. In fact, we have new requirements and challenges in the context of bidirectional model transformation.

Most importantly, we should be sure that a bidirectional model transformation behaves exactly as we want. Such a transformation has more complicated behavior

^{*2}This paper is an extended version of the conference short paper [Hidaka et al., 2011c] presented in ASE'11.

^{*1}Current affiliation is Google

than a unidirectional one. It should be *well-behaved* in the sense that both the forward and backward transformations are consistent with each other and have the roundtrip property [Czarnecki et al., 2009]. As argued in [Stevens, 2007] though, there are semantic issues with many of the existing tools.

Next, bidirectional model transformations should be *compositional* so that they can reuse existing transformations and bigger ones can be constructed from smaller ones. As indicated in the conclusion of [Ehrig et al., 2005], most model transformation languages based on graph transformations are rule-based, describing direct relationships between the source and target models. They are not compositional in the sense that we cannot introduce *intermediate models* for gluing together model transformations. Therefore, rule-based techniques cannot easily support systematic development of model transformations in the large [Klar et al., 2007]. However, composition comes at the cost of efficiency; many unnecessary intermediate models might be produced. Therefore, an optimization method is required to automatically eliminate unnecessary intermediate models during execution.

Furthermore, a bidirectional model transformation should be general enough because it is to be used at various stages of the software development life cycle. It should be able to be applied to different models such as UML diagrams, sequence diagrams, Petri-nets, and even lower level control/data flow graphs. While visual frameworks are useful in high-level design, *general text-based languages* play an important role in developing large-scale transformations, say, to deal with lower level mappings or complex code refactoring. Moreover, we would expect to have a set of language-based tools for type checking (validating) both models and bidirectional model transformations to remove errors before execution, an efficient execution model, and a tool for testing/debugging bidirectional behavior. Apart from a recent attempt in [Wider, 2011] that embed lens [Foster et al., 2007] as a DSL in Scala, which benefits from the type system in the host language, as far as we are aware, no such *language-based modeling environments* have been proposed for bidirectional model transformation.

In this paper, we remedy this situation by proposing a language-based modeling framework called GRoundTram (Graph Roundtrip Transformation for Models), which is carefully designed and implemented for *compositional* development of *well-behaved* and *efficient* bidirectional model transformation at the various stages of software development. Our work is inspired by recent research on bidirectional languages (with well-behavedness) and automatic bidirectionalization in the programming language community [Bohannon et al., 2008, Foster et al., 2007, Hu et al., 2008, Matsuda et al., 2007]. In particular, it has been recently shown [Hidaka et al., 2010] that a graph query algebra UnCAL (Unstructured CALculus) [Buneman et al., 2000], which consists of carefully designed graph constructors, conditional and structural recursion operators, can be fully bidirectionalized. Each graph transformation in UnCAL

has clear bidirectional semantics and is guaranteed to be well-behaved.

This paper is about a successful application of a bidirectional graph query algebra in the programming community to the construction of a framework for developing bidirectional model transformations in the software engineering community. Our main technical contributions are summarized as follows.

Well-Behavedness. We propose a novel *bidirectional graph contraction algorithm* so that we can build well-behaved bidirectional model transformations upon the well-founded bidirectional UnCAL algebra. In fact, there is a gap between UnCAL graphs and the models used in model transformation: graphs in UnCAL are edge-labeled and their equivalence is defined by bisimilarity, while models in model transformation may have labels on both edges and nodes and their equivalence is defined by unique identifiers. We close this gap so that every UnCAL graph has a bidirectional correspondence with a model.

Compositional. We design a user-friendly language UnQL⁺, which is the first *purely functional language* for developing large bidirectional model transformations in a *compositional* way. UnQL⁺ is an extension of the graph query language UnQL [Buneman et al., 2000] with new additional language constructs for graph transformation. We show that any UnQL⁺ program can be correctly translated into an UnCAL construct and inefficiency due to intermediate models in the composition can be automatically eliminated.

Language-based IDE. We implement an integrated development environment GRoundTram, which has a novel tool for validating both models and bidirectional model transformations, an automatic optimization mechanism for improving efficiency, and a powerful debugging environment for testing bidirectional behavior. The system (including sources, documents, and many application examples) is available online [BiG], and being used by research groups besides ourselves for developing some nontrivial applications. Their successes indicate the usefulness of GRoundTram in practice.

This paper is an enhancement of the emerging idea of our previous work that was presented in a short paper [Hidaka et al., 2009b]. Moreover, in relation to the bidirectionalization presented in [Hidaka et al., 2010], a bidirectional contraction stage has been added after stage 2 (epsilon edge and unreachable part elimination stage) of the previous work. With respect to the surface syntax extension presented in [Hidaka et al., 2009a], the present paper can deal with regular path expressions and treat multiple graph databases. We have also added a proof of correctness of a translation to an internal graph algebra and have integrated the verification framework presented in [Inaba et al., 2011] into ours.

The rest of the paper is organized as follows. We begin by demonstrating how the GRoundTram system works in Section 2. We then briefly review the UnCAL bidi-

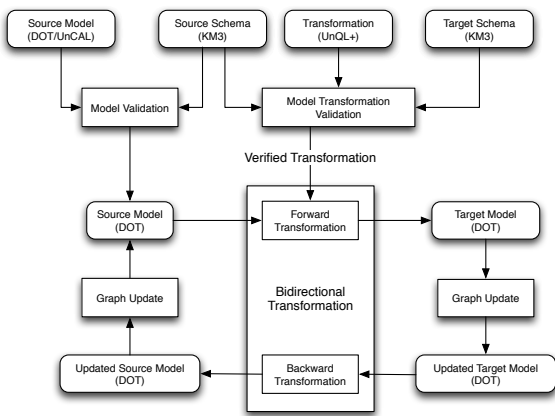


Fig. 1 Overview of GRoundTram




Fig. 2 Snapshot of GRoundTram

rectional framework on which GRoundTram is based in Section 3. After explaining the design architecture of the GRoundTram system, we show in detail the definition of UnQL⁺, the bidirectional graph contraction algorithm, and the translation from UnQL⁺ to UnCAL in Section 4. We evaluate the system in Section 5, discuss the related work in Section 6, and conclude the paper in Section 7.

2 Overview of GRoundTram

Before proceeding with the technical details, let us overview the GRoundTram system to give the reader an idea of what it can do. Figure 1 shows the basic functions the GRoundTram system provides.

2.1 Input

The input to the system is a source model together with its schema, a transformation described in UnQL⁺, and a target model schema. The target model is produced by the forward transformation.

Model. Models are represented by general edge-labeled graphs, which form a general representation of various models. As a running example, consider the class model diagram in Figure 3, which is taken from [Ehrig et al., 2005]. It consists of three classes and two directed associations, and each class has a primary attribute*. This

*For this particular example, the diagram could be further simplified by using simple labeled references instead of model elements for phone number, for example, and nothing prevents us from doing so. However, we believe this example helps us demonstrate the expressiveness of transformations in GRoundTram.

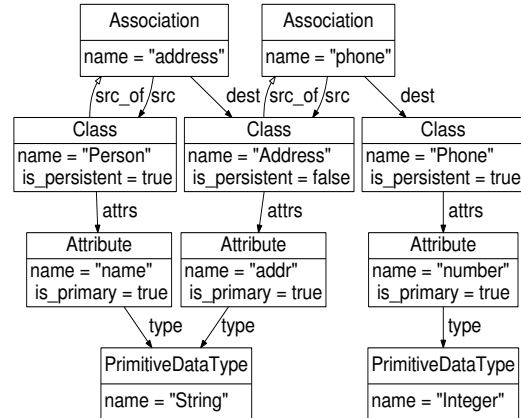


Fig. 3 A Class Diagram




Fig. 4 A Class Model Represented by an Edge-Labeled Graph

model can be represented by the graph in Figure 4, where the information has been moved to the edges. The graph is in the standard DOT format which can be visualized and edited by the popular Graphviz tool [Ellson et al., 2003].

Model Schema (Metamodel). Each model has a structure. For instance, a class diagram has the following structure. A class diagram consists of classes and directed associations between classes. A class is indicated as persistent or non-persistent. It consists of one or more attributes, at least one of which must be marked as constituting the classes' primary key. An attribute type is of a primitive data type (e.g. String, Integer). An association associates classes and is represented here using a model element. KM3 [Jouault and Bézivin, 2006] is used to describe such a model structure [Hidaka et al., 2009a], and its definition can be found in [BiG]. We currently do not support complex OCL constraints in the schema.

Model Transformation. (Forward) Model transformation is described compositionally in UnQL⁺ (Section 4.1), a SQL-like graph query/transformation language. As an example, consider extracting all persistent classes from the class model \$db, and transforming them into

tables by replacing *Attributes* with *Columns*. This can be described compositionally as follows, using the intermediate model $\$persistentClass$.

```

select {tables : $table} where
  $persistentClass in
    (* select classes *)
    (select $class where
      {Association.(src|dest).Class : $class} in $db,
      {is_persistent : {Boolean : true}} in $class),
  $table in
    (* replace Attribute *)
    (replace attrs → $g
      by (select {Column : $a} where
          {attrs.Attribute : $a} in $persistentClass)
        in $persistentClass)

```

2.2 Validation

In order to detect errors during development as early as possible and help users to develop correct models and transformations, GRoundTram provides two validation mechanisms.

Model Validation. The system can verify the conformance of the source and the target model to their associated schemata. In particular, after editing the models, it is important to check that they are in valid states.

Model Transformation Validation. Correct model transformations should always generate a target model conforming to the target schema from any source model satisfying the source schema.

While the model validation is standard, a general model transformation validation is more challenging but more useful for ensuring correct model transformations. As an instance of simple erroneous transformation, suppose the user made an error writing `select $a` instead of `select {Column : $a}` in the previous example. Its outputs do not conform to the schema and hence an error is reported by the system. The check is *automatic* and *static*. Users neither have to provide any test cases by hand, nor execute the transformation for testing; the system automatically finds and displays an example of a source model that reveals the problem (in this case, a class model containing at least one persistent class).

2.3 Bidirectional Transformation

The GRoundTram system is unique in its execution of well-behaved bidirectional transformations, as seen in the lower part of Figure 1.

Forward Transformation. After the user specified the source model and the UnQL⁺ model transformation, the target model is computed by running the transformation with the source model set to the variable $\$db$. Like the source model, the target model can also be exported and be edited in the standard DOT format.

Backward Transformation. The most distinct feature of GRoundTram is the automatic derivation of backward transformations that appropriately propagate modifications on the target models to the source models. There is no need to maintain two separate transformations

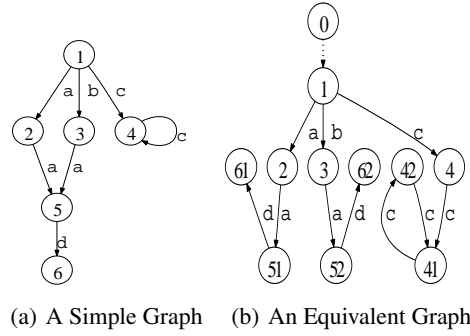


Fig. 5 Graph Equivalence Based on Bisimulation

or to worry about their consistency. Users just write a forward transformation from one model to another in a compositional way, and a corresponding backward transformation is automatically derived.

2.4 Graphical User Interface

The GRoundTram system combines all its functions as an integrated framework with a user-friendly GUI (Figure 2). The user loads a source graph (displayed in the left pane) and a bidirectional transformation written in UnQL⁺. Once they are loaded, the forward transformation can be conducted by pushing the “forward” button (right arrow icon). The target graph appears on the right pane. The user can graphically edit the target graph and apply a backward transformation by pushing the “backward” button (left arrow icon). The source graph can be edited as well, of course. The user can optionally specify the source schema and the target schema and can run the validation by pushing the check button on both panes. The transformation itself can also be checked.

For ease of debugging/understanding the behavior of the bidirectional computation between two models, *trace information* is instantly displayed between the source and target (red part in Figure 2). If subgraphs on either pane are selected, corresponding subgraphs on the other pane are also highlighted. This helps users to understand how a modification on the target affects that on the source, and vice versa.

3 Background: Bidirectional UnCAL

The GRoundTram system is built upon our recent work [Hidaka et al., 2010] on bidirectionalization of UnCAL, a graph algebra known in the database community for graph querying [Buneman et al., 2000]. It has been shown that any unidirectional graph transformation written in UnCAL can be fully bidirectionalized with a backward transformation such that both forward and backward transformations are consistent and well-behaved. We briefly explain the basic results that will be used in this paper.

3.1 Graph Data Model

Graphs in UnCAL are rooted and directed cyclic graphs with no order between outgoing edges. They are edge-labeled in the sense that all information is stored as labels on edges and labels on nodes serve as unique identifiers and have no particular meaning. Figure 5(a)

gives a small example of a directed cyclic graph with six nodes and seven edges. In text, it is represented by

$$\begin{aligned} g &= \{a : \{a : g_1\}, b : \{a : g_1\}, c : g_2\} \\ g_1 &= \{d : \{\}\} \\ g_2 &= \{c : g_2\} \end{aligned}$$

where the notation $\{l_1 : g_1, \dots, l_n : g_n\}$ denotes a set representing a graph which contains n edges with labels l_1, \dots, l_n , each edge pointing to a graph g_i , and the empty set $\{\}$ denotes a graph with a single node. Two graphs g_1 and g_2 can be merged using the set union operation $g_1 \cup g_2$. In addition, the ε -edge is allowed to represent a shortcut between two nodes, and works like the ε -transition used in automata.

Two graphs in UnCAL are considered to be equal if they are *bisimilar*. An intuitive understanding of bisimilarity is that unfolding of cycles and duplication of equivalent subgraphs do not affect the equivalence of graphs, and unreachable parts from the root are ignored. For instance, the graph in Figure 5(b) is equivalent to the graph in Figure 5(a); the new graph has an additional ε -edge (denoted by the dotted line), duplicates the graph rooted at node 5, and unfolds and splits the cycle at node 4.

It is worth noting that bisimilarity plays an important role in bidirectionalization [Hidaka et al., 2010], query optimization [Buneman et al., 2000], and verification of graph transformations [Inaba et al., 2011]. However, bisimilarity is different from the usual equivalence of models whose elements have unique identifiers. We will show how to bridge this gap in Section 4.2.

3.2 UnCAL

The most important feature of UnCAL is that any graph transformation in UnCAL is described by structural recursions or their composition.

A structural recursive function f in UnCAL is a recursive computation scheme on graphs defined by

$$\begin{aligned} f(\{\}) &= \{\} \\ f(\{l : g\}) &= (l, g) \odot f(g) \\ f(g_1 \cup g_2) &= f(g_1) \cup f(g_2) \end{aligned}$$

where \odot is a given binary operator. Different choices of \odot define different recursive functions. For simplicity, the definition above is abbreviated to

$$\mathbf{sfun} \ f(\{l : g\}) = (l, g) \odot f(g).$$

Note that even for a graph g having cycles, the computation of $f(g)$ always terminates under the usual *recursive semantics*, where all recursive calls are memoized and their results are reused to avoid entering infinite loops.

As a simple example, we may use the following recursive function *a2d_xc* to replace all edges labeled a by d and skip all edges labeled c for the graph in Figure 5(a).

$$\mathbf{sfun} \ a2d_xc(\{l : g\}) = \mathbf{if} \ l = \mathbf{a} \ \mathbf{then} \ \{d : a2d_xc(g)\} \\ \mathbf{else} \ \mathbf{if} \ l = \mathbf{c} \ \mathbf{then} \ a2d_xc(g) \\ \mathbf{else} \ \{l : a2d_xc(g)\}$$

We can naturally extend the structural recursion above so that it allows mutual recursion. Any number of mutually recursive functions can be merged into one by using the standard tupling method [Hu et al., 1997].

3.3 Bidirectional Semantics of UnCAL

A query in UnCAL is usually run in the forward direction: under an environment (a mapping from variables to graphs) ρ , a query Q generates a result graph denoted by $\mathcal{F}[[Q]]\rho$.

Let $g = \mathcal{F}[[Q]]\rho$ be a result graph. Assume that a user has edited it into g' . For example, one may add a new subgraph, modify some labels, or delete several edges. In our previous study [Hidaka et al., 2010], we gave *backward semantics* that properly reflect back the editing to the original inputs. Formally speaking, given the modified result graph g' and the original input environment ρ , we presented a method that computes the modified environment $\rho' = \mathcal{B}[[Q]](\rho, g')$.

By “properly reflecting back” (or *well-behaved*), we mean the following two properties hold:

$$\frac{\mathcal{F}[[Q]]\rho = g}{\mathcal{B}[[Q]](\rho, g) = \rho} \text{(GETPUT)}$$

$$\frac{\mathcal{B}[[Q]](\rho, g') = \rho'}{\mathcal{B}[[Q]](\rho, \mathcal{F}[[Q]]\rho') = \rho'} \text{(WPUTGET)}$$

The (GETPUT) property says that if no change is made to the output g , there should be no change in the input environment. The (WPUTGET) property is an unrestricted version of the (PUTGET) property that appeared in [Foster et al., 2005], which requires $g' \in \text{Range}(\mathcal{F}[[Q]])$ and $\mathcal{B}[[Q]](\rho, g') = \rho'$ to imply $\mathcal{F}[[Q]]\rho' = g'$. The (PUTGET) property states that if a result graph is modified to g' , which is in the range of the forward evaluation, this modification can be reflected in the source such that a forward evaluation will produce the same result g' . In contrast, the (WPUTGET) property allows the modified result to be different (this difference is sometimes called the *view side-effect*) from the result obtained by a backward evaluation followed by a forward evaluation, but requires both results to have the same effect on the original source if the backward evaluation is applied again.

Although the (WPUTGET) property is weaker than the (PUTGET) property that forbids view side-effect, this property enables us to make flexible modifications on the result graphs. For example, if the transformation includes a duplication, the target will include two copies of the same data. If a user, being unaware of the duplicates, edits only one of them, then this modification would be forbidden in the (PUTGET) setting, because updating only one of the copy will make the updated target out of the range of the transformation. Instead, we reflect the updates and the user has a chance to do another forward transformation to see the update reflected in another copy in the view.

Even with the flexibility of (WPUTGET) explained above, we reject updating of values that come from transformation, like label d in *a2d_xc* described earlier. Because whatever reflection was made, another forward transformation will always produce d again, so the user’s modification in the target would not have been preserved.

Note that these properties are true under bisimulation equivalence. We additionally need encoding in order to represent models whose equivalence is based on




Fig. 6 GRoundTram Implementation on Bidirectional UnCAL Engine

isomorphism. In our earlier study in [Sasano et al., 2011], we fill this gap by encoding identifiers of the model elements with dedicated edges. In this way, no distinct value-equivalent subgraph encoding different model elements will be contracted.

We use three different editing reflection mechanisms [Hidaka et al., 2010] for different editing operations: edge renaming, edge deletion, and sub-graph insertion. Trace information is used to reflect edge deletion and determine the insertion point in the source. Insertion is handled by using a general inversion technique [Abramov and Glück, 2002] to enumerate the pairs of the updated target and the corresponding source. Unspecified edge label (corresponding to default values) in the source should be filled in by users.

Although we have an SQL-like `select/replace/delete-where` surface syntax (described in section 4.1), bidirectional interpretation of the transformation takes place at the UnCAL level. Therefore, although we can encode join operations by using consecutive `where` clauses, this join operation is not our unit of bidirectionalization (this join operation is translated to nested structural recursions which are units of bidirectionalization), so we can not exploit these high level semantics to reflect changes like relational lenses [Bohannon et al., 2006]. Diskin et al. [2011] pointed out a similar composability issue in the state-based setting under the context of keys. We do not use the concept of a key, so we do not have this problem. We do not have a control over update policy using the notion of key.

Hermann et al. [2011] discuss the correctness of the synchronization algorithm for TGGs. Compared with this related work, our well-behavedness (correctness) reasoning is only at the graph level of representation. We could have better notion of correctness for users at model level as in [Hermann et al., 2011]. These related works are important towards our reasoning about correctness in the model level.

4 Design and Implementation of GRoundTram

Figure 6 depicts the architecture of GRoundTram. We provide a new user-friendly model transformation language UnQL⁺ that is *functional* (rather than rule-

based as in many existing tools) and *compositional* with high modularity for reuse and maintenance, and the architecture handles models that are described by edge-labeled graphs that are general enough to express various models. GRoundTram system runs on the powerful engine of bidirectional UnCAL, which has a set of language-based tools: a bidirectional interpreter [Hidaka et al., 2010], a graph and graph transformation verifier [Inaba et al., 2011], an optimizer to improve efficiency [Hidaka et al., 2011a], and a checker of valid updates in the backward transformation [Nakano et al., 2011]. The key contributions in this implementation are (1) a translation of UnQL⁺ into UnCAL to enable the engine of bidirectional UnCAL to execute UnQL⁺ bidirectionally, and (2) a bidirectional graph contraction algorithm for contracting bisimilar UnCAL graphs so that an ordinary model will have a bidirectional correspondence with an UnCAL graph.

In the rest of this section, we will focus on the exploration of UnQL⁺, the bidirectional graph contraction algorithm, and the translation from UnQL⁺ to UnCAL.

4.1 Model Transformation in UnQL⁺

UnQL⁺ is the language the GRoundTram system provides for users to describe (bidirectional) model transformations. It is an extension of the well-known UnQL [Buneman et al., 2000], a graph querying language, which is compositional and can be implemented by FO (TC) (first order with transitive closure) with PTIME time complexity for graph querying.

Figure 7 gives the core syntax of UnQL⁺. A graph transformation is described by a template expression to construct a graph from graphs that are bound by graph variables. The expression $\{l_1 : t_1, \dots, l_n : t_n\}$ creates a new node having n outgoing edges labeled l_i and pointing to the root of the graph computed from t_i . The union $g_1 \cup g_2$ constructs a graph with a root sharing the roots of g_1 and g_2 . The variable expression $\$g$ returns the graph that is bound by $\$g$ in the environment (i.e., a mapping from variables to graphs). The conditional expression has the usual meaning, i.e., choosing different branch according to the (binding) condition B .

Like other query languages, UnQL⁺ has a convenient template expression `select t where bs` , which is used to select the subgraphs satisfying the a sequence of conditions bs , binds them to variables, and construct a result according to the template expression t . For instance, the following query extracts all persistent classes from the class model in Figure 4, which is assumed to be bound by $\$db$.

```
select $class where
  { Association.(src|dest).Class : $class } in $db,
  { is_persistent : { Boolean : true } } in $class
```

This query returns all bindings of the variable $\$class$ satisfying the two conditions in the *where* clause. The first condition is to find bindings of $\$class$ by matching the *regular path pattern* `Association.(src|dest).Class` with the graph bound by $\$db$, while the second condition is to ensure that the class is persistent.

In model transformations, one often wants to replace

(template)	$T ::= \{L : T, \dots, L : T\} \mid T \cup T \mid \g \mid if BC then T else T \mid select T where B, \dots, B \mid replace $Rp \rightarrow \$G$ by T in T where B, \dots, B \mid extend $Rp \rightarrow \$G$ with T in T where B, \dots, B \mid delete $Rp \rightarrow \$G$ in T where B, \dots, B \mid let sfun $fname(L : \$G) = \dots$ in $fname(T)$
(binding)	$B ::= Gp$ in $\$G \mid BC$
(condition)	$BC ::=$ not $BC \mid BC$ and $BC \mid BC$ or BC \mid isEmpty $(\$G) \mid L = L \mid L \neq L \mid L < L \mid L \leq L$
(label)	$L ::= \$l \mid a$
(label pattern)	$Lp ::= \$l \mid Rp$
(graph pattern)	$Gp ::= \$G \mid \{Lp : Gp, \dots, Lp : Gp\}$
(regular path pattern)	$Rp ::= a \mid _ \mid Rp.Rp \mid (Rp Rp) \mid Rp? \mid Rp* \mid Rp+$

Fig. 7 Syntax of UnQL⁺

a subgraph satisfying a certain condition by another graph, and it is onerous to describe such transformations using *select-where* because some context structure must be copied and propagated. To this end, we introduce three new template expressions, namely, *replace-where*, *extend-where*, and *delete-where*.

- The *replace-where* expression replaces a subgraph with a new graph. For the following *replace-where* expression,

replace $r \rightarrow \$v$ **by** e_1 **in** e_2 **where** b_1, \dots, b_n

the semantics of this expression are that, starting from the root node of e_2 , it traverses every path and replaces the node $\$v$, which is on each path that matches r and satisfies $b_1 \wedge \dots \wedge b_n$, by e_1 . Consider the class model again, prefixing every name of the class by “class_” can be done as follows. Note that “^” is a built-in function for string concatenation.

```
replace  $_*.$ Class.name.string  $\rightarrow$   $\$u$ 
by {"class_"^ $\$name$ } : {} in  $\$db$ 
where { $\$name$  : {}} in  $\$u$ 
```

- The *delete-where* expression is used to describe the deletion of part of the graph.

For the following *delete-where* expression,

delete $r \rightarrow \$v$ **in** e **where** bs

the semantics of this expression are that, starting from the root node of e , it traverses every path and deletes the node $\$v$, which is on each path that matches r and satisfies bs . For instance, we may eliminate all persistent classes by

```
delete Association.(src|dest).Class  $\rightarrow$   $\$c$ 
in  $\$db$ 
where {is_persistent.Boolean : true}
in  $\$c$ 
```

where the subgraph matching $\$c$ will be deleted from its original graph $\$db$.

- The *extend-where* expression describes the extension of a graph with another graph.

For the following *extend-where* expression,

extend $r \rightarrow \$v$ **with** e_1 **in** e_2 **where** bs

the semantics of this expression are that, starting from the root node of e_2 , it traverses every path and extends the node $\$v$, which is on each path that matches r and satisfies bs , with e_1 . For example, we write the following transformation to add date information to each class.

```
extend  $_*.$ class  $\rightarrow$   $\$c$ 
with {date : "2008/8/4"}
in  $\$db$ 
```

These three new template expressions can be automatically translated to structural recursions in UnCAL (see Section 4.3).

Unlike most rule-based model transformation languages, where model transformation composition is not straightforwardly supported [Ehrig et al., 2005], UnQL⁺ is functional and compositional; smaller model transformations can be composed to form bigger ones (see Section 2 and Section 5).

4.2 Bidirectional Graph Contraction

As explained in Section 3, our graph model is based on bisimulation equivalence, which means bisimilar graphs cannot be distinguished. Moreover, since UnCAL is based on bisimulation, a transformation may introduce redundant nodes that are bisimilar to each other. Therefore, after a transformation such redundancy has to be eliminated in a normalization phase.

Fortunately, for any set of graphs that are bisimilar to each other, there exists a unique normal form up to isomorphism and we can obtain the normal form after a transformation by using the partition refinement algorithm of Paige and Tarjan [Paige and Tarjan, 1987]. The algorithm’s complexity is $O(|E| \log |V|)$, where $|E|$ and $|V|$ are the numbers of edges and nodes, respectively, and we consider that this level of complexity would be acceptable in practice. Although this algorithm works on node-labeled graphs, we lift the algorithm to our edge-labeled

graph model by converting edges labeled by l into two unlabeled edges and a node between them labeled by l , as described in [Buneman et al., 2000]. Any bisimilar subgraphs are then contracted to one subgraph, in which no pairs of nodes are bisimilar to each other. In particular, leaf nodes (nodes that have no outgoing edges) are bisimilar to each other, so they all shrink to one node.

We carefully design our contraction algorithm so that it forms a well-behaved bidirectional transformation that has the (GETPUT) and (WPUTGET) properties explained in Section 3.3, in the sense that no modification on the contracted graph results in no modification on the uncontracted graph, while the modified uncontracted graph can be obtained again after re-uncontracting the graph obtained by contracting the modified uncontracted graph. For example, suppose the set of nodes $V_1 = \{v_{1_1}, \dots, v_{1_M}\}$ are contracted to v_1 , and the set of nodes $V_2 = \{v_{2_1}, \dots, v_{2_N}\}$ are contracted to v_2 . Further, suppose an edge (v_1, l, v_2) is inserted in the contracted graph. As a result, edges labeled l are inserted between V_1 and V_2 in the uncontracted graph. If $M > 1$ and $N > 1$, then edges labeled l are inserted only between pairs of nodes that were originally connected in the uncontracted graph, although all-to-all connections from V_1 to V_2 are also well-behaved. If no pair of nodes were originally connected, then the above all-to-all connection is used.

It is worth noting that (WPUTGET) law (instead of PURGET) here is not caused by the duplication of forward transformation. Violation of (PUTGET) law occurs when the modification of the target makes non-bisimilar nodes bisimilar. For example, suppose source graph $\{a:\{b:\{\}\}, a:\{c:\{\}\}\}$. Contraction will produce $\{a:\{b:\{\}\}, a:\{c:\{\}\}\}$. And suppose the label c is modified to b in the contracted graph. Then backward transformation will produce $\{a:\{b:\{\}\}, \{b:\{\}\}\}$. Next forward transformation (contraction) will produce $\{a:\{b:\{\}\}\}$, which is not isomorphic to previous target $\{a:\{b:\{\}\}, \{b:\{\}\}\}$, although they are bisimilar. Since contraction transforms bisimilar graphs to its normal form up to isomorphism, the non-isomorphic targets should be considered different, so only (WPUTGET) law is satisfied.

4.3 Translating UnQL⁺ to UnCAL

UnQL⁺ is different from UnCAL in that it uses four important template expressions, namely *select*, *replace*, *extend*, and *delete*, to describe graph transformations rather than using structural recursion. In this section, we show that all these template expressions can be translated into structural recursions in UnCAL.

The *select* expression, which is inherited from UnQL, can be translated into structural recursion in [Buneman et al., 2000] (the explanation is omitted). According to the semantics described in Section 4.1, the *delete* and *extend* expressions can be defined in terms of the *replace*

expression as:

```
delete Rp → $v in e1 where bs
⇒ replace Rp → $v by {} in e1 where bs
```

```
extend Rp → $v with e1 in e2 where bs
⇒ replace Rp → $v by $v ∪ e1 in e2 where bs
```

Therefore, what we need to show is how the *replace* expression is translated into a structural recursion.

Our idea for this translation is to use structural recursion to simulate the behavior of a deterministic finite automaton (DFA) for finding the nodes in the graph where the *replace* operation is to be applied. For the *select* expression in UnQL, NFA is used to define the structural recursion for finding the nodes in graph to be selected [Buneman et al., 2000]. The reason we use a DFA instead of an NFA is to keep the context correct. The detailed explanation is provided at the end of this section.

Now, consider the following general form of the *replace* expression.

```
replace Rp → $v by e1 in e2 where bs
```

First, we translate the regular path pattern Rp into a DFA $(Q, \Sigma_{\$l}, \delta, q_0, F)$, where $Q = \{q_0, \dots, q_N\}$ is a finite set of states, $\Sigma_{\$l} = \Sigma \cup \{\$l\}$ (where $\Sigma = \{l_0, \dots, l_K\}$) is a finite set of labels used in Rp , $\delta : Q \times \Sigma \rightarrow Q$ is the transition function, $q_0 \in Q$ is the start state, and $F \subseteq Q$ is a set of accept states. We use the special label $\$l$ to denote a label other than those used in Rp .

Next, we introduce $N + 1$ functions h_{q_0}, \dots, h_{q_N} , where h_{q_i} corresponds to state q_i , and define each h_{q_i} as a structural recursion in the following way. For each label $l \in \Sigma_{\$l}$, we define

$$h_{q_i}(\{l : \$v\}) = e_{ij}$$

and construct a graph with e_{ij} by considering two cases. Note that j ranges over transitions from q_i . If $\delta(q_i, l) \notin F$ (i.e., transition from q_i through l does not reach to an accept state), we keep the context by propagating l and continuing the recursive computation by defining


$$e_{ij} = \{l : h_{\delta(q_i, l)}(\$v)\}.$$

Otherwise, we check whether $\$v$ satisfies the condition bs . If it does, we replace the graph with the query result of e_1 satisfying bs :

```
eij = if isEmpty(select {"found"} where bs)
      then {l : h_{\delta(q_i, l)}(\$v)}
      else {l : (select e1 where bs)}.
```

The condition `isEmpty(...)` in the `if`-expression checks whether the condition bs holds. Note that since e_1 might be evaluated to $\{\}$, the checking expression should not be `select e1 where bs`.

By applying the function associated with the initial state to e_2 , we get a UnQL expression having both structural recursions and *select* expressions. Finally, since the *select* expressions can be translated into structural recursions by using the existing method, we can get structural recursions in UnCAL.



```

let sfun  $h_{s0}(\{a : \$v\}) =$  if isEmpty( $e_1$ )
    then  $\{a : h_{s1}(\$v)\}$ 
    else  $\{a : e_2\}$ 
|  $h_{s0}(\{c : \$v\}) =$  if isEmpty( $e_1$ )
    then  $\{c : h_{s1}(\$v)\}$ 
    else  $\{c : e_2\}$ 
|  $h_{s0}(\{\$l : \$v\}) = \{\$l : h_{s0}(\$v)\}$ 
sfun  $h_{s1}(\{a : \$v\}) =$  if isEmpty( $e_1$ )
    then  $\{a : h_{s1}(\$v)\}$ 
    else  $\{a : e_2\}$ 
|  $h_{s1}(\{c : \$v\}) =$  if isEmpty( $e_1$ )
    then  $\{c : h_{s1}(\$v)\}$ 
    else  $\{c : e_2\}$ 
|  $h_{s1}(\{\$l : \$v\}) = \{\$l : h_{s0}(\$v)\}$ 
in  $h_{s0}(\$db)$ 
where  $e_1 \equiv$  select  $\{\text{"found"} : \{\}\}$ 
    where  $\{g : \$u\}$  in  $\$v, \{_{-}.*.d : \$v'\}$  in  $\$v$ 
 $e_2 \equiv$  select  $\$v'$ 
    where  $\{g : \$u\}$  in  $\$v, \{_{-}.*.d : \$v'\}$  in  $\$v$ 
    
```

(c) Translated structural recursion

Fig. 8 A Translation Example




Fig. 9 Source graph and its transformation by replace expression in Figure 8(a)

Example 1. Our algorithm maps the *replace* expression shown in Figure 8(a) to the structural recursion in Figure 8(c) via the DFA obtained from $_{-}.*.(a|c)_{-}$ in Figure 8(b).

This example also demonstrates the user-friendliness of the *replace* syntax, since without this extension, we have to directly code the mutually recursive function in Figure 8(c), after manually constructing the corresponding DFA in Figure 8(b). This manual coding would be error-prone and much more verbose than the *replace* syntax.

Now, let us show the correctness of the above translation from UnQL⁺ to UnCAL. As described in Section 4.1, we interpret the expression

```

replace  $r \rightarrow \$v$  by  $e_1$  in  $e_2$  where  $bs$ 
    
```

so that, starting from the root node of e_2 , it traverses every path and replaces the node $\$v$, which is on each path that simultaneously matches r and satisfies bs , by e_1 (actually **select** e_1 **where** bs should be used instead of e_1 to obtain the bindings from bs). More formally, it is intended to work equivalently to the following pseudo code.

```

let sfun  $f(\{\$l : \$v\}) =$ 
    if MATCH $[r, bs](\$v)$ 
    then  $\{\$l : \text{select } e_1 \text{ where } bs\}$ 
    else  $\{\$l : f(\$v)\}$ 
in  $f(e_2)$ 
    
```

where **MATCH** $[r, bs]$ is a pseudo predicate that is evaluated to be true if and only if $\$v$ is the first node on the path from the root that matches r and satisfies bs . All branches in our desugaring have either the form $\{\$l : \text{select } e_1 \text{ where } bs\}$ or $\{\$l : f(\$v)\}$. Thus, what we have to prove is that the **MATCH** $[r, bs]$ pseudo predicate is correctly encoded in the mutual recursion of structural recursive functions. First, whether the node $\$v$ matches r is encoded in the mutual recursion, which simulates the behavior of the deterministic finite automaton (DFA). This is a standard technique to represent regular patterns. Next, whether the node $\$v$ satisfies bs is encoded in a **if**-expression having **isEmpty**(**select** $\{\text{"found"}\}$ **where** bs) as its condition. This condition holds if and only if bs does not hold, since **isEmpty**(e) holds if and only if e is evaluated into $\{\}$. Thus, this **if**-expression results in **select** e_1 **where** bs when bs holds. Otherwise, it results in keeping the input label as $\$l$, together with calling the function associated with the next state in the DFA. Therefore, the **MATCH** $[r, bs]$ pseudo predicate is correctly encoded in our desugaring code.

Example 2. Figure 9(b) shows the result of using the *replace* expression in Figure 8(a) to transform the cyclic graph in Figure 9(a). This example illustrates what happens if a pattern matches the middle of a cycle. Since the *replace* expression captures the first match along the path, node 4 in Figure 9(a) matches for replacement. Note that DFA in Figure 8(b) only corresponds to a path in the *replace* clause, but the **where** clause also works to

specify the matching node. Therefore, even though node 5 also matches by virtue of the DFA, it does not actually match since it does not satisfy the **where** clause. The obtained binding of $\$v'$ is a subgraph under node 0, bound by the regular path pattern $.*.d$. The subgraph under node 8 in Figure 9(b) is a copy of the subgraph under node 0. The entire correspondence of the source nodes to target nodes is

$$\begin{aligned} 7 &\rightarrow \{0\} \\ 6 &\rightarrow \{10, 7\} \\ 5 &\rightarrow \{9, 6\} \\ 4 &\rightarrow \{5\} \\ 3 &\rightarrow \{4\} \\ 2 &\rightarrow \{3\} \\ 1 &\rightarrow \{2\} \\ 0 &\rightarrow \{8, 1\} \end{aligned}$$

The 'first match' semantics correspond to a translation algorithm which uses **select** to generate subgraphs in the **by** clause, since the **select** construct itself does not conduct recursive replacement as **replace** can.

Necessity of DFA in the translation of UnQL⁺

We give a detailed explanation of why DFA is necessary for translation of **replace/delete/extend** constructs.

Automata created for **replace/delete/extend** play a role that is different from those created for **select**.

For **select**, a regular path pattern (RPP) represents the paths from a node to target nodes reached by the paths and all the subgraphs below the target nodes are unified by graph union \cup and returned as a result. Multiple matches from a given node through identical labels are encoded in an automaton by transitions (nondeterministic branches) to different sub-automata that encodes different subsequent patterns. This is implemented in **sfun** that calls different **sfuns** associated with the target states and unify the graphs returned by those called **sfuns** by graph union \cup . If, during the traversal, no further match is possible for the RPP, then no further transition in the automata is possible (i.e., dead state is reached) and correspondingly the **sfun** returns empty ($\{\}$). Note that there is no semantic problem if the NFA for **select** is determinized. It is just unnecessary to determinize in order to preserve semantics. Also note that by dead states, we mean states with no transitions, not useless states like those unreachable from the initial state.

Instead, **extend** uses RPP to specify target nodes and unify the subgraphs below each of these nodes with a given subgraph by \cup , while the rest of the input graph is kept intact. It is implemented by traversing from the top as **select** does, copying the traversed path, instead of discarding the traversed path in case of **select**. Before reaching the target node during the traversal, if no more match is possible, then the original subgraph should just be returned. This situation is also represented by a special dead state in the automata encoding RPP, and corresponding **sfun** just returns the subgraph. On the other hand, if transitions with an identical label have multiple matches, NFA based encoding would utilize different sub-automata for the targets of these transitions and cor-

responding **sfun** unifies the results by \cup . Suppose input graph $\$db = \{a : \{b : \{x : \{\}\}\}\}$ and transformation

$$\text{extend } (a.b)|(a.c) \rightarrow \$v \text{ with } \{y : \{\}\} \text{ in } \$db$$

The NFA approach would generate an automaton in which target sub-automata from the initial state via transition by **a** would be different, each encoding the rest of the patterns **b** and **c**. So corresponding **sfun** would call two **sfuns** for input **a** and unify them. Then, the former sub-automaton would copy **b** and then produce matching result by unifying $\{x : \{\}\}$ and given graph $\{y : \{\}\}$, thus, $\{b:\{x:\{\},y:\{\}\}\}$. On the other and, the latter sub-automaton would produce non-matching result after copying **b** the original subgraph, thus, $\{b:\{x:\{\}\}\}$. So the unified result would be $\{a:\{b:\{x:\{\},y:\{\}\}\}, a:\{b:\{x:\{\}\}\}\}$. However, semantics of **extend** should return $\{a:\{b:\{x:\{\},y:\{\}\}\}\}$ instead. Although they are trace equivalent, they are not bisimilar. On the contrary, DFA approach produces DFA with only one target state from the transition by label **a** from the initial state. The target sub-automaton is only one. Therefore, corresponding **sfun** produces linear preceding paths before reaching target node of the RPP, i.e., $\{a : \{b : _}\}$. Similar arguments apply for **replace** and **delete**. Therefore, DFA is essential to implement **extend/replace/delete** correctly. It is true that exponential blowup might occur during determinization, however, if the RPP is not complex, i.e., it is not large and corresponding NFA does not include many nondeterministic branches, the number of states will remain moderate. Note that the need of DFA does not come directly from the dead states, but from the problem of unifying the result of subautomaton which might behave differently (some of which might reach dead state and the other reach matching state, as in the example given above).

5 Evaluation and Applications

Here, we demonstrate the power (expressiveness and efficiency) of the GRoundTram system through the development of a known nontrivial (bidirectional) model transformation between UML class diagrams and relational databases models, and highlight its usefulness in practice by giving a list of important applications developed by other groups using it.

5.1 Developing Bidirectional Class2RDB

Class2RDB is a model transformation proposed by Bézivin et al. [2005] as a common benchmark example for all the participants of the Model Transformation in Practice workshop for comparing and contrasting various approaches to model transformation. Class2RDB maps Class models to RDB models. For instance, it can be used to transform the Class model in Figure 3 into the RDB model in Figure 10. Class2RDB maps each persistent class in a Class model to a table in the RDB model. All attributes of the class and its subclasses are mapped to columns in the corresponding table. If a primary attribute belongs to the class, a **pkey** reference from the table model element pointing the corresponding




Fig. 10 An RDB Model

Table 1 Summary of experiments (running time is in CPU seconds)

	direction	no rewriting	rewriting
<i>Class2RDB</i>	forward	1.18	0.68
	backward	14.5	7.90
<i>PIM2PSM</i>	forward	0.08	0.07 (13)
	backward	1.62	0.75
<i>C2Osel</i>	forward	0.04	0.05 (11)
	backward	0.26	0.27
<i>C2Osel'</i>	forward	0.05	0.04 (11)
	backward	2.53	1.26
<i>Ex1</i>	forward	0.036	0.007 (1)
	backward	0.83	0.69

column model element is established. If an attribute belongs to its subclass which is persistent, a reference describing foreign key to the corresponding table model element is established.

We show that $UnQL^+$ is powerful enough to (compositionally) describe the forward transformation (from class diagrams to relational database models), while getting the backward transformation for free in our framework. Figure 11 shows the whole transformation in $UnQL^+$. Let us briefly explain how this $UnQL^+$ program was developed by splitting the transformation into two steps. In the first step (denoted by the binding of `$tables_step1`), every persistent class is mapped to a subgraph representing table, which is connected to subgraphs that correspond to columns originated from attributes of the class and its subclasses. All subclasses are collected by using regular path patterns as shown in Section 4.1. If necessary, references `pkey` and `fkeys` are added by an `extend` construct in $UnQL^+$, provided that the references `refs` of `Fkey` do not point to the referring table because the table may not have been constructed yet. They point to the name of the referring table instead. In the second step (denoted by the binding of `$tables_step2`), each name pointed to by `refs` is replaced by the corresponding table by using a `replace` construct.

5.2 Optimization and Efficiency

Next, we show that the forward and backward transformations can be run efficiently in a scalable manner, while the inefficiency due to composition can be automatically removed through our fusion optimization. The fusion operation was originally designed to merge two successive applications of structural recursions into one [Buneman et al., 2000]. It was later enhanced in [Hidaka et al., 2011a,b].

Table 1 summarizes the performance of bidirectional transformations on various compositional transformations. The test used a MacOSX on a 17 inch MacBookPro, with a 3.06 GHz Intel Core 2 Duo




Fig. 12 Transformation Time v.s. Source Graph Size

CPU. An edge-renaming algorithm was used in the measurement, and no modification was actually made, since it does not significantly affect the running time. The rightmost column shows the running time with rewriting optimization. The number of nodes and edges of the graph (Figure 4) that encoded the Class model (Figure 3) in *Class2RDB* were respectively 70 and 73. The sizes of the source models in the other transformations were similar. (See the analysis at the end of this subsection for the performance with respect to the size of the input graph.) In the table, *PIM2PSM* stands for Platform Independent Model to Platform Specific Model transformation, *C2Osel* for transformation of a customer oriented database into an order oriented database, followed by a simple selection, and *Ex1* for the example from our previous paper [Hidaka et al., 2009a], which in turn was borrowed from [Buneman et al., 2000]. *Ex1* is a composition of two structural recursions. The numbers in parentheses show how often the fusion transformation happened. Our rewrites led to performance improvements in both directions. As the run-time optimization, unreachable parts were removed after every application of the $UnCAL$ structural recursion operator. This run-time optimization is effective when the composed transformation has high selectivity (generates small output from large input), whereas fusion is effective when the selectivity is low. Note that this optimization was turned off for *C2Osel'*. The slowdown in *C2Osel* after rewriting accounts for this trade-off. For the principles of this rewriting optimization, please refer to our papers [Hidaka et al., 2011a,b]. You can test other optimizations, like, e.g., subgraph computation optimization, at our project's website [BiG].

To account for the slowdown of the backward transformation compared with the forward transformation, we took a sample execution profile. The backward transformation decomposes a graph using reachable subgraph extraction computation. In addition, the structural recursion `let sfun ... e_b ... in e_a` involves restoring the input of the argument expression e_a , which in turn requires examining environment produced by the backward transformation of the structural recursion and superposing the resultant graphs in order to restore the entire graph as the updated target of the argument expression. The backward transformation of *PIM2PSM* (after rewriting) uses 15 times more node id comparison operations compared

```

select $tables_step2 where
  $tables_step1 in
    (select $tables where
      {Class:$class} in (select $assoc where {Association.(src|dest):$assoc} in $db),
      {is_persistent.Boolean:true} in $class,
      $dests in (select {Class:$dest} where {(src_of.Association.dest.Class)+:$dest} in $class),
      $related in ((Class:$class) U $dests),
      $cols in (select {cols:{Column:{name:$n,type:$t}}}
        where {Class.attrs.Attribute:{name:$n,type:$t}} in $related),
      $tables in (select {Table:{name:$cname} U $cols} where {name:$cname} in $class),
      $tables in (extend Table -> $table with $pkeys U $fkeys in $tables where
        {cols:$cols} in $table,
        {Column.name.String:{:$cname}} in $cols,
        $pkeys in (select {pkey:$cols} where
          {attrs.Attribute: {is_primary.Boolean:true, name.String:{:$pname:{}}} in $class,
          $cname = $pname),
        $fkeys in (select {fkeys:{Fkey:{cols:$cols, ref:$ref}}} where
          {Class:{is_persistent.Boolean:true,
            attrs.Attribute.name.String:{:$aname:{}}, name:$ref}} in $dests,
          $cname = $aname))),
  $tables_step2 in (replace Table.fkeys.Fkey.ref -> $ref by {Table:$table} in $tables where
    {Table:$table} in $tables_step1,
    {String:{:$rname:{}}} in $ref,
    {name.String:{:$tname:{}}} in $table,
    $tname = $rname)

```

Fig. 11 Class2RDB in UnQL⁺

with the forward transformation. The node comparisons are a result of adding or looking up nodes or edges in sets or maps that are implemented by using balanced binary trees in the OCaml standard library.

Figure 12 shows how the size of the source model affects time to execute the *a2d_xc* transformation described in Section 3, in both directions. Lattice-like regularly shaped strongly connected graphs are used as the source. These execution times match the complexity of PTIME mentioned in Section 4.1 for relatively large (several thousand nodes) graphs.

5.3 Other Applications

The GRoundTram website [BiG] has a bunch of examples, big and small, and all the examples presented in this paper can be tried through the demo website. In addition, we would like to give the reader a rough idea about the current status of GRoundTram uses by listing applications that have been or are being developed by other groups: they include bidirectional feature model transformation [Wang et al., 2010](Peking University), bidirectional transformation between VDM (Vienna Development Method) specifications and Java implementation [Ishikawa, 2011] (another group at National Institute of Informatics), bidirectional transformation between Simulink diagrams and UML diagrams [Takayuki Kozawa, 2011] (Waseda University), bidirectionalizing ATL (ATLAS Transformation Language) with GRoundTram [Sasano et al., 2011] (Shibaura Institute of Technology), and co-evolution of Java models and codes [Yu et al., 2012] (Open University & Shanghai Jiao Tong University). Moreover, Chen [2010] at Shanghai Jiao Tong University identifies GRoundTram as a potential means to synchronize the behavior model of concurrent systems. All these activities indicate the promise of GRoundTram as a practical tool.

As an example of our own experience in [Yu et al., 2012], we extensively used the UnQL⁺ syntax as an internal representation that was automatically generated by

the blinkit[†] tool. Although it was not directly used by users, it was a concise way of describing the user editing operations to Java code, and also for exploiting high-level optimization opportunities at this syntax level.

6 Related Work

Besides the related work in the introduction, we highlight some others related to graph-based model transformation and linguistic approach to bidirectional transformation.

Our work is much related to research on model transformation based on graph transformation. AGG [Taentzer, 2003] is a rule-based visual tool that supports typed (attributed) graph transformations, including type inheritance and multiplicities. Triple Graph Grammars (TGG) [Giese and Wagner, 2006, Königs and Schürr, 2006] is intended as a declarative specification of model-to-model integration rules. MOFLON [Amelunxen et al., 2006] implements TGG and adopts the visual notation of QVT Relational, the OMG standard bidirectional model transformation language. Giese and Hildebrandt [2009] proposed a model synchronization algorithm based on TGG that can synchronize large-scale models. Guerra et al. [2011] proposed more general notion called "inter-modelling", where a specification can be compiled into different operational mechanisms not only for model-to-model transformation but also for model matching and model traceability. To perform forward and backward transformations, a pattern specification is compiled into Triple Graph Grammar (TGG) operational rules.

Different from these rule-based approaches, ours is a functional one that supports model transformation composition and its automatic optimization. As far as we are aware, this is the first nontrivial *functional and algebraic framework for model transformation*.

[†]<http://computing-research.open.ac.uk/linkit/>

This study was inspired by the recent linguistic approach to bidirectionalization of the tree transformation [Bohannon et al., 2008, Foster et al., 2005, 2007, Hu et al., 2008, Matsuda et al., 2007] for tree data synchronization. One important feature of this approach is clear bidirectional semantics, something that is missing from most of the existing bidirectional model transformation systems [Stevens, 2007]. Although some attempts at using the linguistic approach have been made [Antkiewicz and Czarnecki, 2007, Xiong et al., 2007], it remains a challenge to provide a general bidirectional framework for graphs which are more complicated than trees, and our work is a big step in this direction.

GRoundTram has grown out of our two-year effort to realize the emergent idea presented in our short paper [Hidaka et al., 2009b]. UnQL⁺ is based on the graph query language UnQL [Buneman et al., 2000] but it is significantly extended with a powerful language construct *replace* that can handle transformation context. It is also worth noting that a simple *replace* expression was studied in [Hidaka et al., 2009a] but it can neither deal with regular path expressions nor treat multiple graph databases.

7 Conclusions

We proposed a novel algebraic framework called GRoundTram to support systematic development of bidirectional model transformation. Different from many rule-based frameworks, our framework is functional and algebraic, and based on a graph algebra and structural recursion. Our new framework supports systematic development of model transformations in a compositional manner, has a clear semantics for bidirectional model transformation, and can be efficiently implemented.

This is our first step towards *bidirectional model programming*, a linguistic framework to support systematic development of model transformation programs. In the future, we will look more into relation between the rule-based approach and the algebraic and functional approach, and see how to integrate them into a more powerful framework for bidirectional model transformation. One initial attempt has already been made by integrating GRoundTram with ATL [Sasano et al., 2011], and we plan to continue this line of research by collaborating with the AtlanMod [Atl] team. In such an integration of model transformation frameworks, we need to automate the graph encoding of UML-like generic diagrams that are not yet fully automated. Several attempts towards automation had been made [Zhu et al., 2012], and we plan to deal with Eclipse Modeling Framework (EMF) models as well. Another direction towards integration is the approach by Wider [2011]. In this approach, asymmetric lens was implemented as an internal DSL in Scala. Although it is not graph-based, the Java-friendliness of the host language may make it easier to integrate with the model driven engineering framework.

Last but not least, we currently do not have any explicit

control over update policy in backward transformations. The only property we have is well-behavedness described by the (GETPUT) and the (WPUTGET). Even if multiple source models may be possible as a result of backward transformation, programmer do not have a way to choose them. Controlling the choice is one of our important future work.

Acknowledgments

We would like to thank the anonymous reviewers for their thorough comments and constructive suggestions to improve the paper. The anonymous ASE'11 reviewers for the prior version of this paper also made helpful suggestions. We also thank Kazutaka Matsuda, Kazuyuki Asada, and Isano Sasano for their valuable discussions with us. The research was supported in part by the Grand-Challenging Project on the "Linguistic Foundation for Bidirectional Model Transformation" of the National Institute of Informatics, and a Grant-in-Aid for Scientific Research for Encouragement of Young Scientists (B) No. 20700035, and a Grant-in-Aid for Scientific Research (B) No. 22300012.

References

- The AtlanMod team web site. <http://www.emn.fr/x-info/atlanmod/>.
- The BiG project web site. <http://www.biglab.org/>.
- Sergei M. Abramov and Robert Glück. Principles of inverse computation and the universal resolving algorithm. In *The Essence of Computation*, pages 269–295, 2002.
- C. Amelunxen, A. Königs, T. Röttschke, and A. Schürr. MOFLON: A Standard-Compliant Metamodeling Framework with Graph Transformations. In A. Rensink and J. Warmer, editors, *Model Driven Architecture - Foundations and Applications: Second European Conference*, volume 4066 of *Lecture Notes in Computer Science (LNCS)*, pages 361–375, Heidelberg, 2006. Springer Verlag, Springer Verlag.
- Michał Antkiewicz and Krzysztof Czarnecki. Framework-specific modeling languages with round-trip engineering. In *MoDELS 2006: Proceedings of the 9th International Conference on Model Driven Engineering Languages and Systems*, pages 692–706. Springer-Verlag, 2006.
- Michał Antkiewicz and Krzysztof Czarnecki. Design space of heterogeneous synchronization. In *GTTSE '07: Proceedings of the 2nd Summer School on Generative and Transformational Techniques in Software Engineering*, 2007.
- Jean Bézivin, Bernhard Rumpe, Andy Schürr, and Laurence Tratt. Model transformation in practice workshop announcement. In *Satellite Events at the MoDELS 2005 Conference*, pages 120–127. Springer-Verlag, 2005.

- Aaron Bohannon, Benjamin C. Pierce, and Jeffrey A. Vaughan. Relational lenses: a language for updatable views. In Stijn Vansumneren, editor, *PODS*, pages 338–347. ACM, 2006.
- Aaron Bohannon, J. Nathan Foster, Benjamin C. Pierce, Alexandre Pilkiewicz, and Alan Schmitt. Boomerang: resourceful lenses for string data. In George C. Necula and Philip Wadler, editors, *POPL '08: ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages*, pages 407–419. ACM, 2008.
- Peter Buneman, Mary F. Fernandez, and Dan Suciu. UnQL: a query language and algebra for semistructured data based on structural recursion. *VLDB Journal: Very Large Data Bases*, 9(1):76–110, 2000.
- Yuting Chen. A bidirectional graph transformation approach to analysis of concurrent software models. In *Software Engineering and Service Sciences (ICSESS), 2010 IEEE International Conference on*, pages 339 – 343, july 2010. doi: 10.1109/ICSESS.2010.5552447.
- Krzysztof Czarnecki, J. Nathan Foster, Zhenjiang Hu, Ralf Lämmel, Andy Schürr, and James F. Terwilliger. Bidirectional transformations: A cross-discipline perspective. In *International Conference on Model Transformation (ICMT 2009)*, pages 260–283. LNCS 5563, Springer, 2009.
- Zinovy Diskin, Yingfei Xiong, and Krzysztof Czarnecki. From state- to delta-based bidirectional model transformations: the asymmetric case. *Journal of Object Technology*, 10:6: 1–25, 2011.
- Hartmut Ehrig, Karsten Ehrig, Claudia Ermel, Frank Hermann, and Gabriele Taentzer. Information preserving bidirectional model transformations. In *Proceedings of the 10th international conference on Fundamental approaches to software engineering, FASE'07*, pages 72–86, Berlin, Heidelberg, 2007. Springer-Verlag. ISBN 978-3-540-71288-6.
- Karsten Ehrig, Esther Guerra, Juan de Lara, Laszlo Lengyel, Tihamér Levendovszky, Ulrike Prange, Gabriele Taentzer, Dániel Varró, and Szilvia Varró-Gyapay. Model transformation by graph transformation: A comparative study. Presented at *MTiP 2005*. <http://www.inf.mit.bme.hu/FTSRG/Publications/varro/2005/mtip05.pdf>, 2005.
- John Ellson, Emden R. Gansner, Eleftherios Koutsofios, Stephen C. North, and Gordon Woodhull. Graphviz and dynagraph - static and dynamic graph drawing tools. In *GRAPH DRAWING SOFTWARE*, pages 127–148. Springer-Verlag, 2003.
- J. Nathan Foster, Michael B. Greenwald, Jonathan T. Moore, Benjamin C. Pierce, and Alan 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.
- J. Nathan Foster, Michael B. Greenwald, Jonathan T. Moore, Benjamin C. Pierce, and Alan Schmitt. Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem. *ACM Trans. Program. Lang. Syst.*, 29(3), 2007.
- Miguel Garcia. Bidirectional synchronization of multiple views of software models. In *Proceedings of DSML-2008*, volume 324 of *CEUR-WS*, pages 7–19, 2008.
- Holger Giese and Stephan Hildebrandt. Efficient model synchronization of large-scale models. Technical Report 28, Hasso Plattner Institute at the University of Potsdam, 2009.
- Holger Giese and Robert Wagner. Incremental model synchronization with triple graph grammars. In *MoDELS 2006: Proceedings of the 9th International Conference on Model Driven Engineering Languages and Systems*, pages 543–557. Springer Verlag, 2006.
- John Grundy, John Hosking, and Warwick B. Mugridge. Inconsistency management for multiple-view software development environments. *IEEE Trans. Softw. Eng.*, 24(11):960–981, 1998.
- Esther Guerra, Juan Lara, and Fernando Orejas. Intermodelling with patterns. *Software & Systems Modeling*, pages 1–30, 2011. ISSN 1619-1366. doi: 10.1007/s10270-011-0192-1.
- Frank Hermann, Hartmut Ehrig, Fernando Orejas, Krzysztof Czarnecki, Zinovy Diskin, and Yingfei Xiong. Correctness of model synchronization based on triple graph grammars. In *Lecture Notes in Computer Science*, volume 6981, pages 668–682. Springer, 2011. ISBN 978-3-642-24484-1.
- Soichiro Hidaka, Zhenjiang Hu, Hiroyuki Kato, and Keisuke Nakano. Towards a compositional approach to model transformation for software development. In *SAC'09: Proceedings of the 2009 ACM symposium on Applied Computing*, pages 468–475, New York, NY, USA, 2009a. ACM.
- Soichiro Hidaka, Zhenjiang Hu, Hiroyuki Kato, and Keisuke Nakano. A compositional approach to bidirectional model transformation. In *ICSE Companion*, pages 235–238. IEEE, 2009b.
- Soichiro Hidaka, Zhenjiang Hu, Kazuhiro Inaba, Hiroyuki Kato, Kazutaka Matsuda, and Keisuke Nakano. Bidirectionalizing graph transformations. In *ACM SIGPLAN International Conference on Functional Programming*, pages 205–216. ACM, 2010.
- Soichiro Hidaka, Zhenjiang Hu, Kazuhiro Inaba, Hiroyuki Kato, Kazutaka Matsuda, Keisuke Nakano, and Isao Sasano. Marker-directed Optimization of UnCAL

- Graph Transformations. In *Proceedings of 21st International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR 2011)*, LNCS, pages 123–138, Odense, Denmark, 2011a.
- Soichiro Hidaka, Zhenjiang Hu, Kazuhiro Inaba, Hiroyuki Kato, Kazutaka Matsuda, Keisuke Nakano, and Isao Sasano. Marker-directed Optimization of UNCAL Graph Transformations (revised version). Technical Report GRACE-TR-2011-06, GRACE Center, National Institute of Informatics, November 2011b.
- Soichiro Hidaka, Zhenjiang Hu, Kazuhiro Inaba, Hiroyuki Kato, and Keisuke Nakano. GRoundTram: An integrated framework for developing well-behaved bidirectional model transformations (short paper). In *26th IEEE/ACM International Conference On Automated Software Engineering*, pages 480–483. IEEE, 2011c.
- Zhenjiang Hu, Hideya Iwasaki, Masato Takeichi, and Akihiko Takano. Tupling calculation eliminates multiple data traversals. In *ACM SIGPLAN International Conference on Functional Programming (ICFP'97)*, pages 164–175, Amsterdam, The Netherlands, June 1997. ACM Press.
- Zhenjiang Hu, Shin-Cheng Mu, and Masato Takeichi. A programmable editor for developing structured documents based on bidirectional transformations. *Higher-Order and Symbolic Computation*, 21(1-2):89–118, 2008.
- Kazuhiro Inaba, Soichiro Hidaka, Zhenjiang Hu, Hiroyuki Kato, and Keisuke Nakano. Graph-transformation verification using monadic second-order logic. In Peter Schneider-Kamp and Michael Hanus, editors, *PPDP*, pages 17–28. ACM, 2011.
- Fuyuki Ishikawa. Towards customizable and bidirectionally traceable transformation between vdm++ and java. In *The 9th Overture/VDM Workshop*, June 2011.
- Frédéric Jouault and Jean Bézivin. KM3: A DSL for metamodel specification. In *Formal Methods for Open Object-Based Distributed Systems*, pages 171–185. LNCS 4037, Springer, 2006.
- Felix Klar, Alexander Königs, and Andy Schürr. Model transformation in the large. In Ivica Crnkovic and Antonia Bertolino, editors, *ESEC/SIGSOFT FSE*, pages 285–294. ACM, 2007.
- Alexander Königs and Andy Schürr. Tool integration with triple graph grammars - a survey. *Electronic Notes in Theoretical Computer Science*, 148(1):113–150, February 2006.
- Ralf Lämmel. Coupled Software Transformations (Extended Abstract). In *First International Workshop on Software Evolution Transformations*, November 2004.
- Kazutaka Matsuda, Zhenjiang Hu, Keisuke Nakano, Makoto Hamana, and Masato Takeichi. Bidirectionalization transformation based on automatic derivation of view complement functions. In *12th ACM SIGPLAN International Conference on Functional Programming (ICFP 2007)*, pages 47–58. ACM Press, October 2007.
- Keisuke Nakano, Soichiro Hidaka, Zhenjiang Hu, Kazuhiro Inaba, and Hiroyuki Kato. Simulation-based graph schema for view updatability checking of graph queries. Technical Report GRACE-TR11-01, GRACE Center, National Institute of Informatics, May 2011.
- Robert Paige and Robert Tarjan. Three partition refinement algorithms. *SIAM Journal of Computing*, 16(6):973–988, 1987. doi: <http://dx.doi.org/10.1137/0216062>.
- Isao Sasano, Zhenjiang Hu, Soichiro Hidaka, Kazuhiro Inaba, Hiroyuki Kato, and Keisuke Nakano. Toward bidirectionalization of ATL with GRoundTram. In *Theory and Practice of Model Transformations, Fourth International Conference, ICMT 2011*, volume 6707 of LNCS, pages 138–151. Springer, June 2011.
- Perdita Stevens. Bidirectional model transformations in QVT: Semantic issues and open questions. In Gregor Engels, Bill Opdyke, Douglas C. Schmidt, and Frank Weil, editors, *Proc. 10th MoDELS*, volume 4735 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2007.
- Perdita Stevens. A landscape of bidirectional model transformations. In Ralf Lämmel, Joost Visser, and João Saraiva, editors, *Generative and Transformational Techniques in Software Engineering II*, pages 408–424. Springer-Verlag, Berlin, Heidelberg, 2008. ISBN 978-3-540-88642-6.
- Perdita Stevens. Bidirectional model transformations in qvt: semantic issues and open questions. *Software and System Modeling*, 9(1):7–20, 2010.
- Gabriele Taentzer. AGG: A graph transformation environment for modeling and validation of software. In John L. Pfaltz, Manfred Nagl, and Boris Böhlen, editors, *AGTIVE*, volume 3062 of LNCS, pages 446–453. Springer, 2003.
- Takayuki Kozawa. Bidirectional transformation with UML model for Simulink model maintainability improvement (in Japanese). <http://www.washi.cs.waseda.ac.jp/papers/2011/submission/lw070119.pdf>, February 2011. Summary of the bachelor's thesis at the Department of Computer Science, Waseda University.
- Bo Wang, Zhenjiang Hu, Qiang Sun, Haiyan Zhao, Yingfei Xiong, and Hong Mei. Supporting feature model refinement with updatable view. Technical Report GRACE-TR-2010-05, GRACE Center, National Institute of Informatics, May 2010.

Arif Wider. Towards combinators for bidirectional model transformations in scala. In Anthony M. Sloane and Uwe Aßmann, editors, *SLE*, volume 6940 of *Lecture Notes in Computer Science*, pages 367–377. Springer, 2011. ISBN 978-3-642-28829-6.

Yingfei Xiong, Dongxi Liu, Zhenjiang Hu, Haiyan Zhao, Masato Takeichi, and Hong Mei. Towards automatic model synchronization from model transformations. In *22nd IEEE/ACM International Conference on Automated Software Engineering (ASE 2007)*, pages 164–173. ACM Press, November 2007.

Yijun Yu, Yu Lin, Zhenjiang Hu, Soichiro Hidaka, Hiroyuki Kato, and Lionel Montrieux. blinkit: Maintaining Invariant Traceability through Bidirectional Transformations. In *Proceedings of the 34th International Conference on Software Engineering (ICSE 2012)*, pages 540–550, Zurich, Switzerland, June 2012.

Yiqing Zhu, Tao Zan, Soichiro Hidaka, and Zhenjiang Hu. iGRT: A generic interface for GRoundTram. Technical Report GRACE-TR-2012-06, GRACE Center, National Institute of Informatics, 2012.