

Semantics of OCL Specified with QVT*

Slaviša Marković and Thomas Baar

École Polytechnique Fédérale de Lausanne (EPFL)
School of Computer and Communication Sciences
CH-1015 Lausanne, Switzerland
e-mail: {slavisa.markovic, thomas.baar}@epfl.ch

The date of receipt and acceptance will be inserted by the editor

Abstract The Object Constraint Language (OCL) has been for many years formalized both in its syntax and semantics in the language standard. While the official definition of OCL's syntax is already widely accepted and strictly supported by most OCL tools, there is no such agreement on OCL's semantics, yet. In this paper, we propose an approach based on meta-modeling and model transformations for formalizing the semantics of OCL. Similarly to OCL's official semantics, our semantics formalizes the semantic domain of OCL, i.e. the possible values to which OCL expressions can evaluate, by a metamodel. Contrary to OCL's official semantics, the evaluation of OCL expressions is formalized in our approach by model transformations written in QVT. Thanks to the chosen format, our semantics definition for OCL can be automatically transformed into a tool, which evaluates OCL expressions in a given context.

Our work on the formalization of OCL's semantics resulted also in the identification and better understanding of important semantic concepts, on which OCL relies. These insights are of great help when OCL has to be tailored as a constraint language of a given DSL. We show on an example, how the semantics of OCL has to be redefined in order to become a constraint language in a database domain.

Key words QVT, OCL Semantics, Graph-transformations, DSL

* This work was supported by Swiss National Scientific Research Fund under the reference number 200020-109492/1.

1 Introduction

The OCL is proven to be a very versatile constraint language that can be used for different purposes in different domains, e.g., for restricting meta-model instances [1], for defining UML profiles [2], for specifying business rules [3], for querying models [4, 5] or databases [6].

Due to the lack of parsers, OCL was used in its early days often in an informal and sketchy style, what had serious and negative consequences as Bauerdick et al. have shown in [7]. Nowadays, a user can choose among many OCL parsers (e.g. OSLO [8], Eclipse Model Development Tool (MDT) for OCL [9], Dresden OCL Toolkit [10], Octopus [11], Use [12], OCLE [13]), which strictly implement the abstract syntax of OCL defined in the OCL standard [14].

The situation is less satisfactory when it comes to the support of OCL's semantics by current OCL tools. While most of the tools now offer some kind of evaluation of OCL expressions in a given system state, none of the tools is fully compliant with the semantics defined in the OCL standard. We believe that the lack of semantic support in OCL tools is due to the lack of a clear and implementation-friendly specification of OCL's semantics. Interestingly, the *normative* semantics of OCL¹ given in the language standard [14], Section 10: *Semantics Described using UML* is also formalized in form of a metamodel, but, so far, this metamodel seems to be poorly adopted by tool builders.

In this paper we present a new approach for formulating a metamodel-based semantics of OCL. Defining a semantics for OCL basically means (1) to define the so-called semantic domain, in which OCL expressions are evaluated, and (2) to specify the evaluation process for OCL expressions in a given context.

The semantic domain for OCL is given by all possible system states. Since a system state can be visualized by an object diagram, the semantic domain is (almost) defined by the official UML metamodel for object diagrams. There are two major problems to be solved when defining the semantic domain based on the definition of object diagrams. Firstly, UML's metamodel for object diagrams does not define the semantics of OCL's predefined types, such as *Integer*, *Real*, *String*, *Set(T)*, etc. However, this problem has been already recognized in the OCL standard and an additional package (named *Values*) for the OCL metamodel has been proposed. We will, to a great extent, reuse the *Values* package in our approach. Secondly, the metamodel for object diagrams implicitly assumes the existence of solely one object diagram at any moment of time. This becomes a major obstacle as soon as more than one system state is relevant for the definition of OCL's semantics (and this is really the case when defining the semantics of OCL's post-conditions). We propose for this problem a solution which

¹ There is also an *informative semantics* given in Annex A of [14], which is formulated in a set-theoretical style and goes back to the dissertation of M. Richters [15].

is fundamentally different from the one chosen in the normative semantics and which leads, as we think, to a much simpler metamodel for the semantic domain of OCL.

The *evaluation of OCL expressions* is specified in our approach by model transformations, which are in turn described as QVT rules [16]. In order to improve readability, we use in this paper a visualization of QVT rules, which is inspired from graph-grammars. All QVT rules presented in this paper are also available in its textual form. The complete set of rules can be downloaded, together with all relevant metamodels, from [17]. Note that the QVT rules are executable on QVT-compliant engines, what is demonstrated by our OCL tool ROCLET[18], which uses internally the QVT rules for the evaluation of OCL expressions. The QVT engine of our choice was Together Architect 2006 [19] which offers mature support for QVT editing, and debugging.

To summarize, our semantics for OCL has the following characteristics:

- The semantics is directly executable. Contrary to a paper-and-pencil semantics, OCL developers can immediately see by using a tool (e.g. ROCLET), how the semantics applies in a concrete scenario.

To our knowledge, only the semantics of OCL given by Brucker and Wolff ([20,21]) has the same characteristics and can be executed in the OCL tool HOL-OCL.

- The semantics is defined on top of the official metamodels for OCL's abstract syntax and UML class- and object-diagrams. Consequently, the semantic definition becomes an integral part of the already existing language definitions for UML and OCL.

However, we had to redefine some of the existing metamodels due to some obvious inconsistencies, which would have prevented us from completely implementing our approach.

- The target audience for our semantics are developers, who use OCL in practice. No familiarity with mathematical and logical formalisms is presumed. In order to understand the semantics, only some knowledge of metamodeling and QVT is required.
- The semantics is presented in a modular way. This allows to easily define, starting from our semantics of OCL, the semantics of another constraint language, which is tailored to a given Domain-Specific Language (DSL). Similarly, one could also create a new dialect for OCL in the context of UML; for example, one could decide to abandon OCL's concept of being a three-valued logic and to allow only two Boolean values *true* and *false*.

The last point highlights the flexibility of our approach. This flexibility is an important step forward to the vision originally formulated by the PUML group (see, e.g., [22]) to treat OCL not just as one monolithic language but rather as a family of languages, which can be applied in many different domains and can adapt easily to different requirements from these domain while still sharing a substantial amount of common semantic concepts, libraries, etc.

This paper is a revised and enhanced version of [23]. While [23] concentrates on the evaluation semantics for invariants, we have added to this paper also rules for the evaluation of pre-/postconditions. Furthermore, many rules were redesigned with the aim to make OCL's underlying semantic concepts more explicit and to make evaluation rules more reusable in other language definitions. We also added a section on tailoring the semantics of OCL towards the needs of a DSL.

The rest of the paper is organized as follows. In Sect. 2, we sketch our approach and show, by way of illustration, a concrete application scenario for our semantics. The basic evaluation steps are formalized by QVT rules in Sect. 3. The formalized QVT rules have to be consistent to each other, at least to a certain degree. Achieving consistency is more likely, if the underlying semantic concepts are made more explicit. Section 4 proposes a list of semantic concepts and discusses their impact on evaluation rules. Section 5 shows the flexibility of our approach and presents a stepwise adaptation of OCL's semantics, so that the adapted version can be used as a constraint language for a given DSL. In Sect. 6, we compare our approach with existing approaches for formalizing the semantics of constraint languages. Section 7 draws conclusions.

2 A Metamodel-Based Approach for OCL Evaluation

In this section we briefly review the technique and concepts our approach relies on and illustrate with a simple example the evaluation of OCL constraints. We concentrate on the evaluation of an invariant constraint in a given state. The difficulties arising from the evaluation of pre-/postconditions are described in Sect. 4.

2.1 Official Metamodels for UML/OCL

We base our semantics for OCL on the official metamodels for UML and OCL. We support the last finalized version of OCL 2.0 [14]. However, since our approach had the requirement to be integrated in the OCL tool RO-CLET, which currently does only support UML1.5 diagrams, we refer also in this paper to UML1.5 as the metamodel of the UML part, on which OCL constraints rely. Figures 1 and 2 show the parts of the UML and OCL metamodels that are relevant for this paper. Please note that Fig. 1 contains also in its upper part a metamodel of the semantic domain of class diagrams.

2.2 Changes in the OCL Metamodel

In order to realize our approach in a clear and readable way, we had to add a few metaassociations, -classes, and -attributes to the *Values* package part

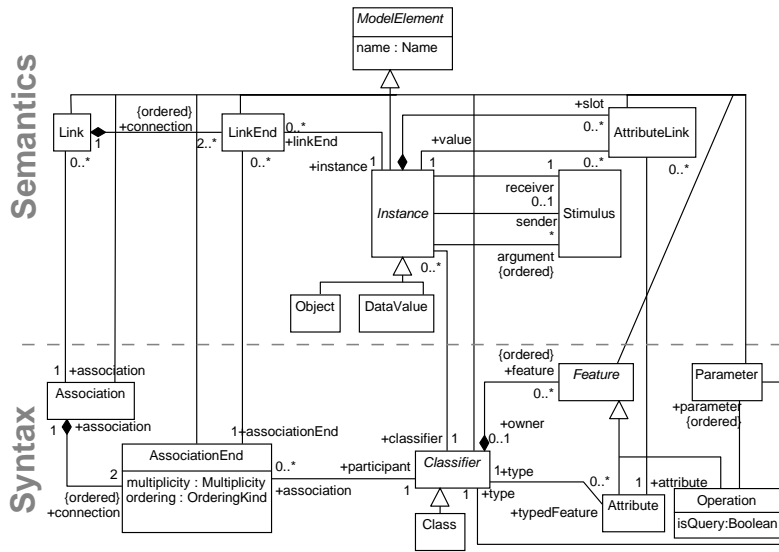


Fig. 1 Metamodel for Class Diagrams - Syntax and Semantics

of the official OCL metamodel (see Fig. 3). The metaclass *OclExpression* has a new association to *Instance*, what represents the evaluation of the expression in a given object diagram. We revised slightly the concepts of bindings (association between *OclExpression* and *NameValueBinding*) and added to class *LoopExp* two associations *current* and *intermediateResult*, and one attribute *freshBinding*. Furthermore, the classes *StringValue*, *IntegerValue*, etc. have now attributes *stringValue*, *integerValue*, etc. what makes it possible to clearly distinguish a datatype object from its value. We have created two new metaclasses *StateTransition* and *ObjectMap* that are used in evaluations of pre and post-conditions. Metaclass *ObjectMap* has two metaassociations with metaclass *Instance* and is used to relate two *Instances* in a pre and a post-state. Metaclass *StateTransition* has two metaassociations with *Stimulus* representing an Operation that corresponds to a given *StateTransition* or a sent message. Stimulus itself is used to keep the track about an operation invocation: receiver and sender of a message, and operation arguments.

2.3 Evaluation

We motivate our approach to define OCL's semantics with a small example. In Fig. 4, a simple class diagram and one of its possible snapshots is shown. The model consists of one class **Stock** with two attributes: **capacity** and **numOfItems**, both of type *Integer*, representing capacity of **Stock** and the current number of items it has, respectively. The additional constraint attached to the class **Stock** requires that the current number of items in a




Fig. 2 Metamodel for OCL - Syntax

stock must always be smaller than the capacity. The snapshot shown in the right part of Fig. 4 satisfies the attached invariant because for each instance of `Stock` (class `Stock` has only one instance in the snapshot) the value of `numOfItems` is less than the value of attribute `capacity`. In other words, the constraint attached to the class `Stock` is evaluated on object `s` to `true`.

In order to show how the evaluation of an OCL constraint is actually performed on a given snapshot, we present in Fig. 5 the simplified state of the Abstract Syntax Tree as it is manipulated by an OCL evaluator. Step (a)-(b) performs the evaluation of the leaf nodes. Depending on the results of these evaluations, step (b)-(c) performs evaluation of nodes at the middle level. Finally, the last step (c)-(d) performs evaluation of the top-level of the AST. Please note that in this example we were not concerned about concrete binding of the variable `self`. The problem of variable binding is discussed in Sect. 2.4.

The basic idea of our approach is that an OCL constraint can be analogously evaluated by annotating directly the OCL metamodel instance instead of the AST.

Figure 6 shows the instance of the OCL metamodel representing the invariant from Fig. 4. Here, we stipulate that all expressions have not been evaluated yet because for each expression the link `val` to metaclass `Instance`




Fig. 3 Changed Metamodel for OCL - Semantics




Fig. 4 Example - Class Diagram and Snapshot




Fig. 5 Evaluation of OCL expressions seen as an AST: (a) Initial AST (b) Leaf nodes evaluated (c) Middle nodes evaluated (d) Complete AST evaluated

is missing. Please note that here we assume that in all expressions variable **self** is bound to the object *o*. For the sake of readability this information is omitted in Figures 6 and 7.

The final state of the metamodel instance, i.e. after the last evaluation step has been finished, is shown in Fig. 7. What has been added compared to the initial state (Fig. 6) is highlighted by thick lines. The evaluation of the top-expression (*OperationCallExp*) is a *BooleanValue* with *booleanValue* attribute set to *true*, the two *AttributeCallExpressions* are evaluated to two

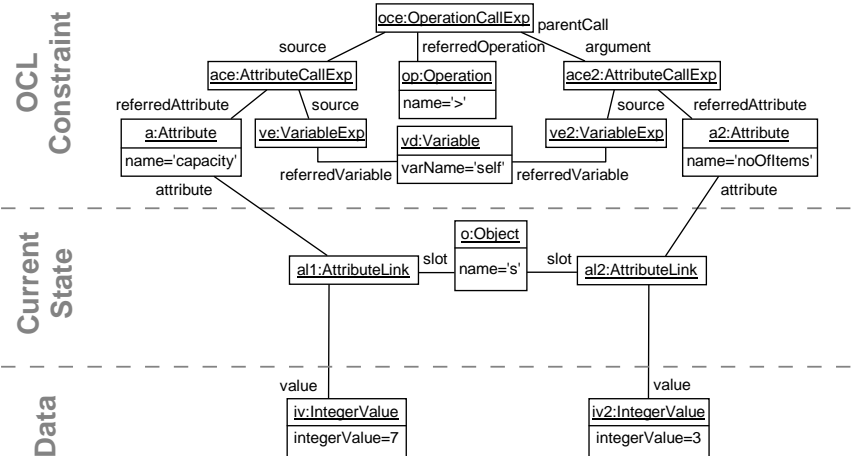


Fig. 6 OCL Constraint Before Evaluation

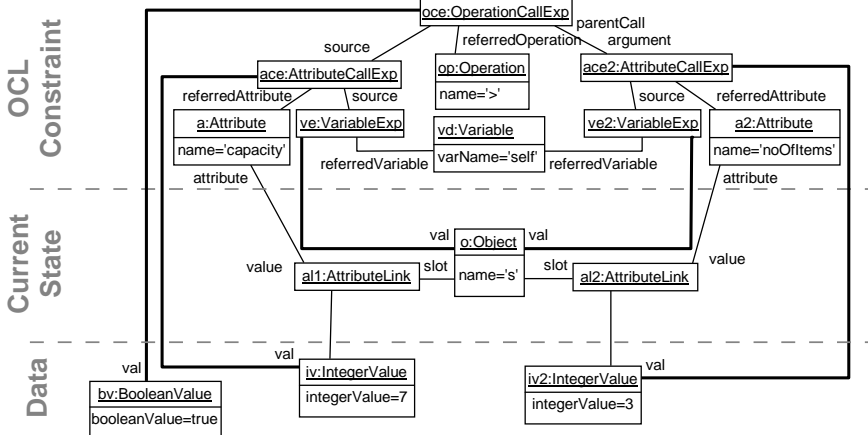


Fig. 7 OCL Constraint After Evaluation in a Given Snapshot

IntegerValues with values 7 and 3, and each *VariableExp* is evaluated to *Object* with name *s*.

2.4 Binding

The evaluation of one OCL expression depends not only on the current system state on which the evaluation is performed but also on the binding of free variables to current values. The binding of variables is realized in the OCL metamodel by the class *NameValueBinding*, which maps one free variable name to one value. Every OCL expression can have arbitrarily many bindings, the only restriction is the uniqueness of variable names within the set of linked *NameValueBinding* instances.

In the invariant of the **Stock** example we have used one free variable, called *self*. Although *self* is a predefined variable in OCL, it can be treated the same way as all other variables, which are introduced in *LoopExp*. For example, the invariant

```
context Stock inv :
  self.capacity > self.numOfItems
```

can be rewritten as

```
Stock.allInstances->forAll( self |
  self.capacity > self.numOfItems)
```

The binding of variables is done in a top-down approach. In other words, variable bindings are passed from an expression to all its sub-expressions. Some expressions do not only pass the current bindings, but also add/change bindings. An example for adding new value-name bindings will be presented in more details in Sect. 3 where the evaluation rules for *iterate* and *let* expressions are explained.

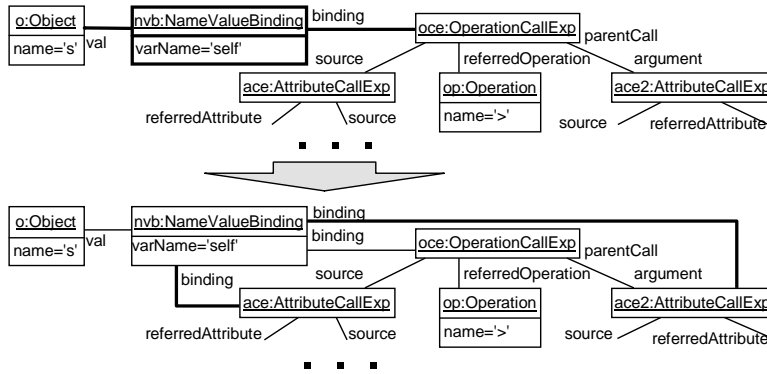


Fig. 8 Binding Passing

Figure 8 shows the process of binding passing on a concrete example. In the upper part, the initial situation is given: The top-expression already has one binding *nvb* for variable *self*. In the lower part of the figure, all sub-expressions of the top-expression are bound to the same *NameValueBinding* as the top-expression.

3 Core Evaluation Rules Formalized as Model Transformations

The previous section has shown the main idea of our approach: we annotate the evaluation result of each (sub)expression directly to the corresponding instance of class *OclExpression* in the OCL metamodel. What has not been specified yet are the evaluation steps themselves, for example, that an *AttributeCallExp* is always evaluated to the attribute value on that object to

which the source expression of *AttributeCallExp* evaluates. As shown below, these evaluation steps will be formally given in form of model transformation rules.

Although the graph-transformation rules are generally readable and understandable nicely, their number can become quite high if one wants to accommodate all peculiarities of OCL (e.g. undefined values, flattening of collections, *@pre* in postconditions, etc.). In order to structure the semantics definition, we will present in this section the core version of evaluation rules for certain types of expressions and will explain in the next Section 4 how this core rules have to be extended/adapted in order to reflect all semantic concepts of OCL.

3.1 Model Transformation Rules

For the specification of evaluation rules we use the formalism of model transformations, more precisely a graphical syntax QVT (Query/View/Transformation) rules [16].

For our application scenario of QVT rules, source and target model are always instances of the same metamodel; the metamodel for UML/OCL including the small changes we have proposed in Sect. 2. Each QVT rule consists of two patterns (LHS, RHS), which are (incomplete) instantiations of the UML/OCL metamodel. When a QVT rule is applied on a given source model, a LHS matching sub model of the source model is searched. Then, the target model is obtained by rewriting the matching sub model by a new sub model that is derived from RHS under the same matching. If more than one QVT rule match on a given source model, one of them is non-deterministically applied. The model transformation terminates as soon as none of the QVT rules is applicable on the current model.

While in the conference version of this paper [23] we have stucked to the official syntax of QVT rules, we take now the freedom to introduce some additional shorthand notations which will help to improve the conciseness and the readability of evaluation rules. One source of the complexity of the rules given in [23] is, that the LHS contains two sub-patterns; one for the structure to look for the current constraint under evaluation (e.g. an *AttributeCallExp*) and one for the structure in the state, in which the constraint is evaluated. The RHS has again two patterns; one for the updated structure of the OCL constraint and one for the structure in the state. Since the evaluation of OCL expressions does not have side-effects on the state in which the expression is evaluated, the state-subpattern of LHS must be the same as the subpattern of RHS.

In order to avoid the redundancy of having the same subpatterns in LHS in RHS, our evaluation rules contain besides LHS and RHS a third part called *Context*, that specify the structures in the input, which must be available when applying the rule but which are not changed (see Fig. 9 for a comparison of the old and the new form of evaluation rules). The




Fig. 9 Format of evaluation rules used in [23] and in this paper

Context part is optional. For the core rules presented in this section, the *Context* will encode the assumed structures in the current state, in which the OCL expression is being evaluated. When it comes to the evaluation of pre-/postconditions, we will see in the next section that the *Context* can also contain even more information. Besides the structures that describe the system state, *Context* can also contain an optional part with data values that are necessary for the evaluation of the rules.

3.2 Binding Passing

Before the source expression can be evaluated, the current binding of variables has to be passed from the parent expression to all its subexpressions. Figure 10 shows the transformation rule for *OperationCallExp*. When applying this rule, the binding of the parent object `oce` (represented by a link from `oce` to the multiobject `nvb` in LHS) is passed to subexpressions `oe` and `aoe` (links from `oe` and `aoe` to `nvb` are established in RHS). Analogous rules exist for all other kinds of OCL expressions which have subexpressions. For the (subclasses of) *LoopExp* (see below) one needs also additional rules for handling the binding because the subexpressions are evaluated under a different binding than the parent expression.

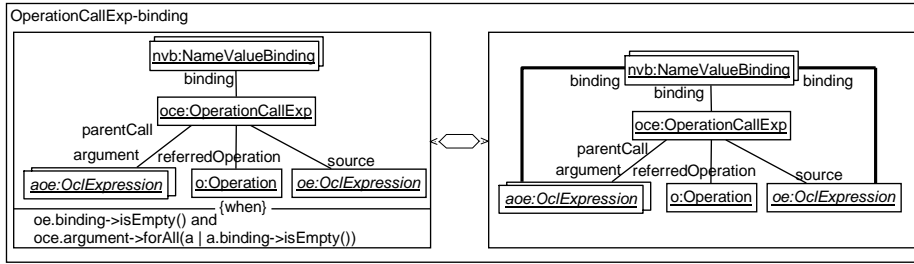


Fig. 10 Binding of an expression

3.3 A Catalog of Core Rules

Each OCL expression is an instance of the metaclass *OclExpression* in the OCL metamodel; more precisely – since *OclExpression* is an abstract metaclass – an instance of one of the non-abstract subclasses of *OclExpression*. For each of these non-abstract metaclasses, the semantics definition must have at least one evaluation rule.

The semantics of a constraint language such as OCL can be split along this syntactic dimension (in Section 4, we will see that it is useful to have also another dimension for the semantics). However, it is not always appropriate to organize a catalog of evaluation rules based on the metaclasses from the abstract syntax metamodel. Sometimes, evaluation rules for different metaclasses are very similar so that these evaluation rules could be put into the same category (for example, *Navigation Expressions*). But there is also the opposite case, where instances of the same metaclass are evaluated using very different mechanisms, what is a sign for a wrong granularity of metaclasses in the metamodel (for example, *OperationCallExp*).

We propose to organize the evaluation rules for OCL based on *Navigation Expressions*, *Operation Expressions*, *Loop Expressions*, *Variable Expressions*, *Literal Expressions*, *If Expressions*, *Message Expressions*², *Let Expressions*, *State Expressions*³, *Tuple Expressions*. For the class of *Operation Expressions*, it is useful to distinguish expressions that refer (1) to predefined operations from the OCL library, (2) to queries defined by the user in the underlying class model, (3) to typecheck or cast operations (e.g. *oclIsTypeOf()*), (4) to *allInstances()*.

Here, we discuss only the most representative rules. The main goal is to demonstrate that the evaluation of all kinds of OCL expressions can be formulated using graph-transformations in an intuitive but precise way.

3.3.1 Navigation Expressions OCL expressions of this category are, for example, instances of *AttributeCallExp* and *AssociationEndCallExp*. Such

² *Message Expressions* can occur only in post-conditions and are ignored here.

³ We consider as the semantic domain of our evaluation only object diagrams in which the objects do not have a reference to an explicit state given in a state diagram. Consequently, *State Expressions* are ignored here.

expressions are evaluated by 'navigating' from the object, to which the source expression is evaluated, to that element in the object diagram, which is referenced by the attribute or association end.

AttributeCallExp The semantics of *AttributeCallExp* is specified by the rule *AttributeCallExp-evaluation* given in Fig. 11. The evaluation of *ace* is datavalue *d*, which is also the value of the attribute *a* for object *o*. Note, that we stipulate in the LHS, that *oc*, the source expression of *ace*, has been already evaluated to object *o*.

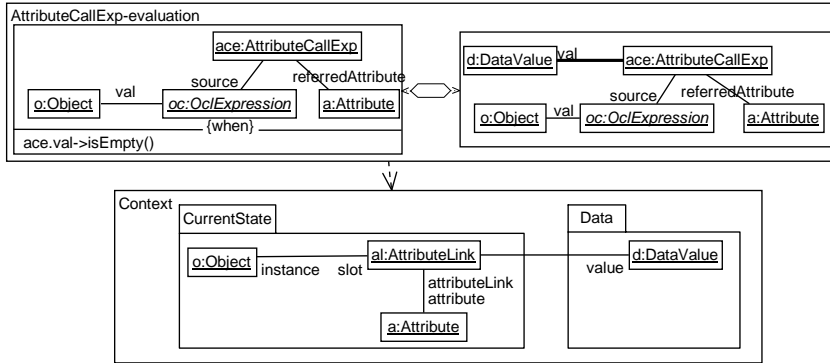


Fig. 11 Attribute Call Expression Evaluation

AssociationEndCallExp We discuss here only the case of navigating over an unordered association end with multiplicity greater than 1 (the case of multiplicities equal to 1 is very similar to *AttributeCallExp*). The rule shown in Fig. 12 specifies that the value of *aece* is a newly created object of type *SetTypeValue* whose elements refer to all objects *o2* that can be reached from object *o* via a link for *ae*. Again, object *o* is the evaluation of source expression *oe*. The rule shown in Fig. 12 contains at few locations the multiplicities 1-1 at the link between two multiobjects, for example at the link between *le2* and *l*. This is an enrichment of the official QVT semantics on links between two multiobjects. Standard QVT semantics assumes that a link between two multiobject means that each object from the first multiobject is linked to every object from the second multiobject, and vice versa. This semantics is not appropriate for the situation shown in Fig. 12 where each element of multiobject 1 must be connected only to one element from multiobject *le2*, and vice versa. By using 1-1 multiplicities, we indicate a non-standard semantics of links between two multiobjects.

3.3.2 Operation Expressions

Expressions Referring to Predefined Operations Expressions from this category are instances of the metaclass *OperationCallExp* but the called operation is a predefined one, such as $+$, $=$. These operations are declared and

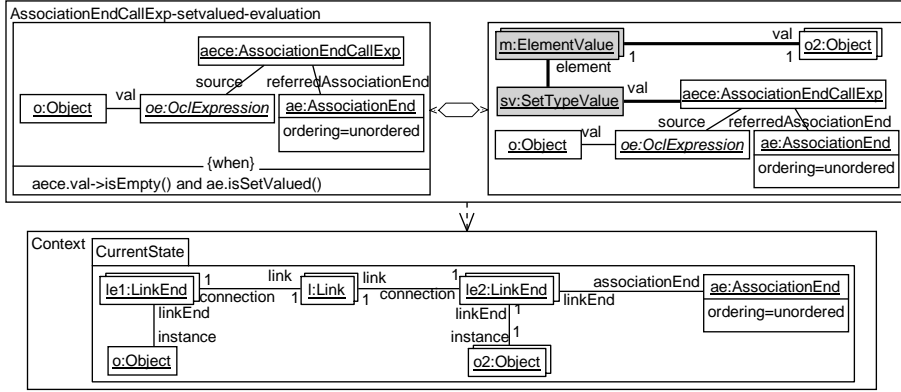


Fig. 12 Association End Call Expression Evaluation that Results in Set of Objects

informally explained in the chapter on the OCL library in [14]. As an example, we explain in the following the semantics of operation “=” (equals). We show only two rules here, one specifies the evaluation of equations between two objects, and the other the evaluation of equations between two integers.

In Fig. 13, the evaluation is shown for the case that both subexpressions `oe1`, `oe2` are evaluated to two objects `o1` and `o2`, respectively. In this case, the result of the evaluation is `bv` of type *BooleanValue* with attribute *booleanValue* `b`, which is *true* if the evaluations of `oe1` and `oe2` are the same object, and *false* otherwise.

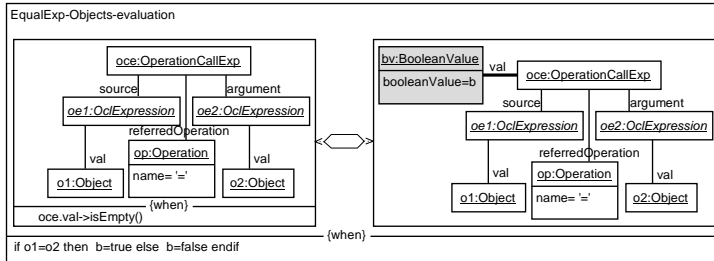


Fig. 13 Equal Operation Evaluation for Objects

If `oe1` and `oe2` evaluate to *IntegerValue*, the second QVT rule shown in Fig. 14 is applicable and the result of evaluation will be an instance of *BooleanValue* with attribute *booleanValue* set to *true* if the attribute *integerValue* of `iv1` is equal to *integerValue* of `iv2`, and to *false* otherwise.

Expressions Referring to a User-defined Query If a user-defined query is used in an OCL constraint, then the semantics of the used query must be specified by a body-clause, which is attached to the query. The query might also have attached a pre-condition, which must evaluate to *true* in the




Fig. 14 Equal Operation Evaluation for Integers

current situation. Otherwise, the query-expression is evaluated to *undefined*. If the pre-condition evaluates to *true*, then the value of the *OperationCallExp* is the same as the evaluation of the clause body under the current argument binding.

Fig. 15 shows evaluation rules for user-defined queries. The first rule creates a set of *NameValueBindings* for the expressions in *precondition* and *body*. Every *NameValueBinding* from this set corresponds to exactly one argument of the *OperationCallExp* *opce*. The second rule performs evaluation of the query in such a way, that if the *precondition* does not evaluate to *true* then the result of the evaluation will be *undefined*, otherwise the result of evaluation of the *body*. One problem, however, is, that the body-expression might contain again an *OperationCallExp* referring to *op*, i.e. the definition of *op* is recursive. Recursive query definitions lead in some but not all cases to infinite loops during the evaluation. Brucker et al. propose in [24] that recursive query definitions should be checked by the user for unfounded recursions, but this would require substantial analysis effort.

Expressions for Typecheck and Typecast To this group belong all *OperationCallExps* referring to the predefined operation *oclAsType* and *oclIsTypeOf*, *oclIsKindOf*. The operation *oclAsType* makes a cast of the source expression to the type specified in the argument. If this cast is successful, the whole expression is evaluated to the same object as the source expression. If the cast is not successful (i.e., the evaluation of the source expression is an object whose type does not conform to the type given in the argument), then the whole expression is evaluated to *undefined*. Because we treat the evaluation to *undefined* in the next Section 4 in a general manner, we skip the rule for *oclAsType* here.

The rules for *oclIsTypeOf* and *oclIsKindOf* are very similar; Fig. 16 shows the rule for *oclIsKindOf*.

allInstances()-Expressions The predefined operation *allInstances()* yields all existing objects of the specified type and all its subtypes. The rule is shown in Fig. 17. Note that the multiobject *os* represents according to the QVT semantics the maximal set of objects *o*, for which the condition given in the *when*-clause of the *Context* holds.

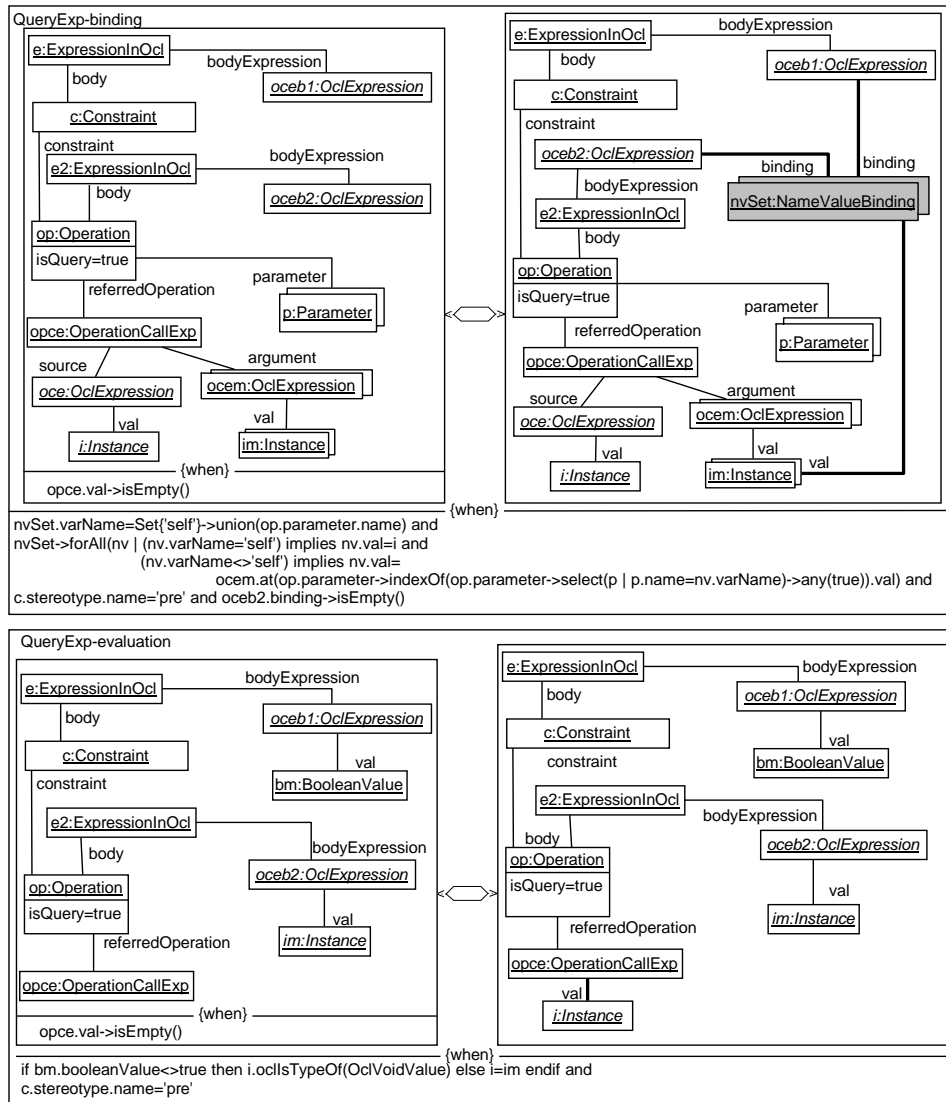

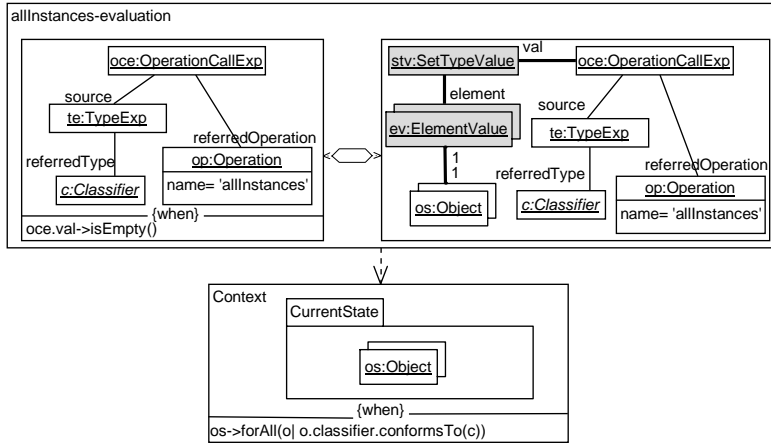


Fig. 15 Evaluation of an expression referring to a query

3.3.3 Loop Expressions Iterator expressions are those in OCL which have as the main operator one from *select*, *reject*, *forAll*, *iterate*, *exists*, *collect*, *any*, *one*, *collectNested*, *sortedBy*, or *isUnique*. Since all these expressions can be expressed by macros based on *iterate*, it is sufficient to refer for their semantics just to the semantics of *iterate*.

In the Fig. 18 are shown evaluation rules that describe the semantics of *iterate*.

Fig. 16 Evaluation Rule for `oclIsKindOf`Fig. 17 Evaluation Rule for `allInstances`

The rule *Iterate-Initialisation* makes a copy of evaluation of the source expression, and assigns it under the role *current* to *ie*. Furthermore, one *NameValueBinding* is created and assigned to the body expression. The name of the *NameValueBinding* is the same as the name of *result* variable, and its value is the same as the value of the *initExpression* for the *result* variable. For some technical reasons, the attribute *freshBinding* of *ie* is set to *false*.

The rule *Iterate-IteratorBinding* updates the binding on body expression *oe* for the iterator variable *v* with a new value *vp*. The element with the same value *vp* is chosen from the collection *current* and is removed afterwards from this collection. The attribute *freshBinding* is set to *true* and the binding for *oe* has changed.




Fig. 18 Iterate - Evaluation Rules

The rule *Iterate-IntermediateEvaluation* updates the binding for the variable with the same name as the result variable of `ie` based on the new evaluation of `oe`. Furthermore, the value of attribute `freshBinding` is flipped and the evaluation of body expression `oe` is removed.

The final rule *Iterate-evaluation* covers the case when the collection `current` of `ie` is empty. In this case the value of `ie` is set to that value which is bound to the `NameValueBinding` with the same name as the result variable.

3.3.4 Variable Expressions Figure 19 shows the evaluation rule for *VariableExp*. When this rule is applied, a new link is created between *Variable-*

Exp and the value to which *NameValueBinding*, with the same name as *VariableDeclaration*, is connected.




Fig. 19 Variable Expression Evaluation

3.3.5 Literal Expressions In Fig. 20, the evaluation of *IntegerLiteralExp* is shown. By applying this rule, a new *IntegerValue* is created whose attribute *integerValue* has the same value as the attribute *integerSymbol* for expression *ie*. Note, that this type of expressions does not need variable bindings because their evaluation does not depend on the evaluation of any variable.




Fig. 20 Integer Literal Expression Evaluation

3.3.6 If-Expressions Figure 21 shows the evaluation rule for an *if* expression. The result of the evaluation depends on the value to which condition expression *c* is already evaluated. As it is stated in the *when* clause of the rule, if the value of the condition is *true* then the result of the evaluation will be the value of the *thenExpression*, otherwise it will be value of the *elseExpression*. Please note that in this example we don't deal with evaluation to undefined and that this aspect of OCL will be discussed later.

3.3.7 Let-Expressions The evaluation of *Let-Expressions* is a little bit different from the other rules because it changes *NameValueBinding* for its subexpressions (similarly to *LoopExpression*). The evaluation rules for *LetExp* are shown in Fig. 22. The first rule performs binding of the *Let-variable* to the value to which *initExpression* evaluates (by creating a new *NameValueBinding* instance), and then passes this *NameValueBinding* to the *in* part of the expression. The second part specifies that result of evaluation of an *LetExp* will be the same as evaluation of its *in* expression.

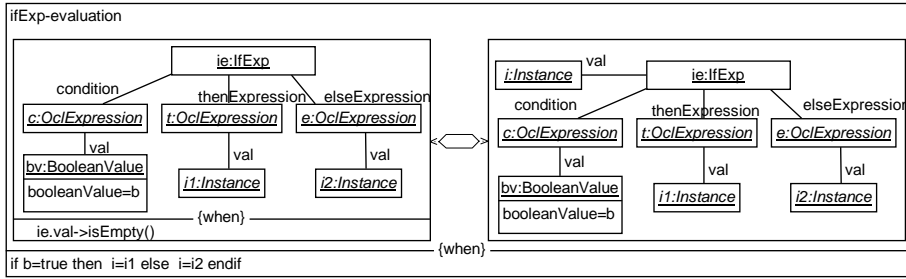


Fig. 21 If-Expression Evaluation


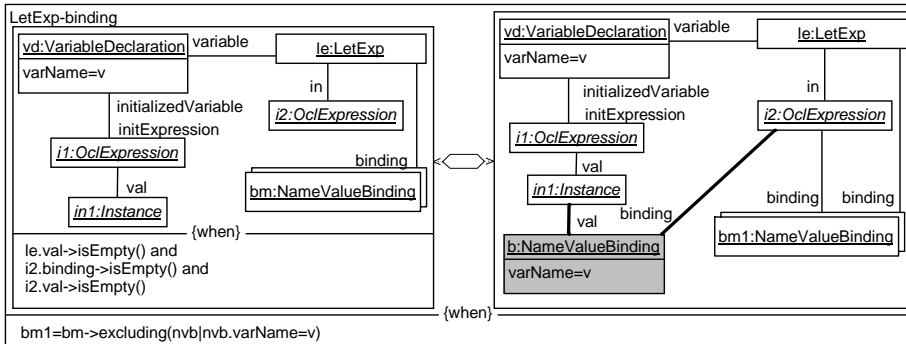


Fig. 22 Let Expression: Binding and Evaluation

3.3.8 Tuple Expressions In Figure 23, the evaluation rule for *TupleExp* is shown. This rule consists of three parts. The first part creates a temporary *TupleValue* object that will become the result of evaluation once all *TupleLiteralParts* are traversed. The middle rule shows the core semantics of *TupleExpression* evaluation. This rule will be executed as many times as there are *TupleLiteralParts* in the expression. Each time this rule is triggered, a new *AttributeLink* is created and attached to the temporary *TupleValue*. This newly created *AttributeLink* will point to one attribute from the tuple type, and to the value that *TupleLiteralPart* has. The third rule is used to create the final value of the *TupleExp*.

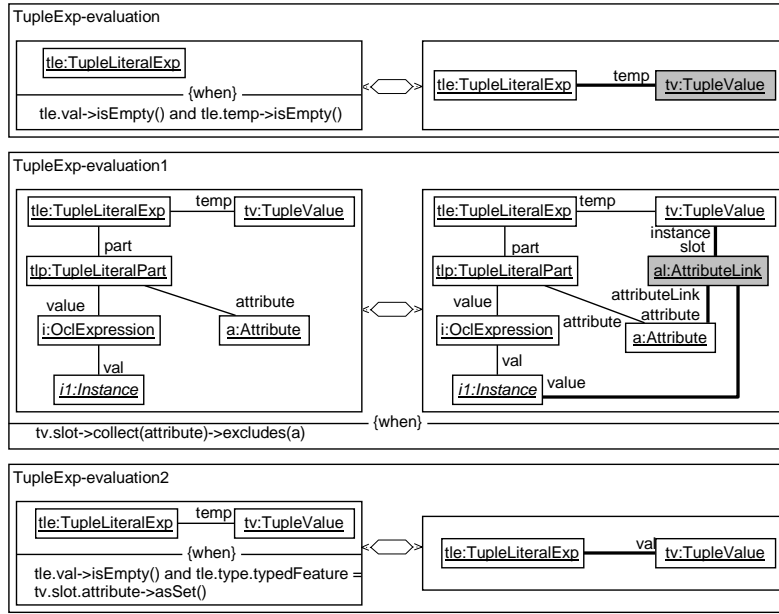


Fig. 23 Tuple Expression Evaluation

3.4 Syntactic Sugar

Many pre-defined OCL operations are defined as an abbreviation for more complex terms. For instance, the operation `exists` can be simulated by the operation `iterate`. More precisely, expressions of form

$$\text{coll} \rightarrow \text{exists}(x \mid \text{body}(x))$$

can be rewritten to

$$\text{coll} \rightarrow \text{iterate}(x; \text{acc} : \mathbf{Boolean} = \text{false} \mid \text{acc} \mathbf{or} \text{body}(x))$$

This rewriting step can also be expressed as a graph-transformation rule what would make the rule for evaluating the pre-defined operation `exists` superfluous.

Figure 24 shows a QVT rule that transforms one `exists` expression into corresponding `IterateExp`. RHS of the rule states that a new `IterateExp` is created, new `VariableDeclaration`, new `BooleanLiteralExp` with `booleanSymbol false`. The source of the expression and the iterator remain the same as for the `exists` operation. The body expression is modified and after the transformation it represents the disjunction of the previous body and the newly created variable expression that refers to the new `VariableDeclaration`. In the `when`-clause, we state an additional constraint that the `varName s` used in the newly created `VariableDeclaration` is not yet used as a name by any of the free variables in the body. Note that this constraint was not specified for the textual representation of the transformation and that would mean that body expression does not contain any free variable with the name `acc`.

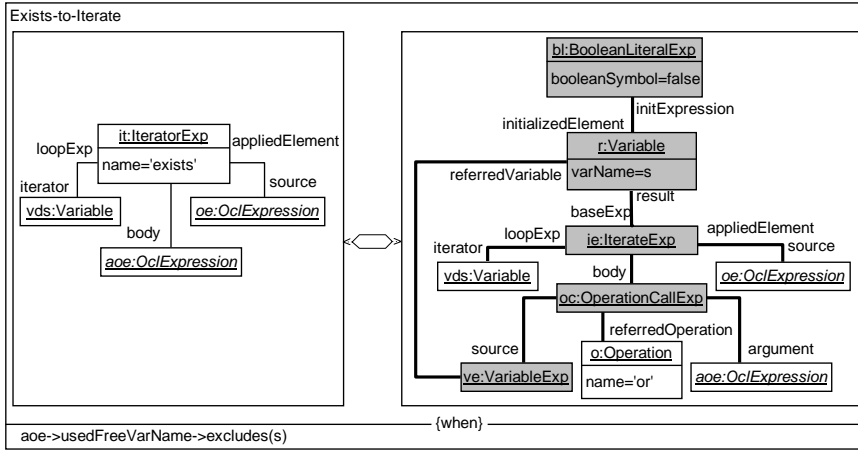


Fig. 24 Transforming Exists expression to an iterate expression

4 Semantic Concepts in OCL

In the previous section, the most important evaluation rules for each of the possible kinds of OCL expressions were given. The rules basically describe the necessary evaluation steps in a given state, but they do not reflect yet the complete semantics of OCL. For example, nothing has been said yet on how an operation contract consisting of pre-/postconditions is evaluated, how to handle the *@pre* construct in postconditions, under which circumstances an expression is undefined, etc. These are examples for *additional semantic concepts*, that are realized in OCL but which are most likely not realized in every other constraint language. Besides the syntactic dimension already explained in Sect. 3.3 for the categorization of rules, the additional semantic concepts form a second dimension for the rule categorization. We have identified the following list of semantic concepts, which must be taken into account when formulating the final version of evaluation rules (note that in Sect. 3.3 only the rudimentary version of evaluation rules has been shown).

- evaluation of operation contracts (pre-/postconditions)
- evaluation to undefined (including strict evaluation with respect to undefined, with some exceptions)
- dynamic binding when invoking a query
- non-deterministic constructs (*any()*, *asSequence()*)⁴

In the next subsections, we discuss the semantical concepts that have the most impact on the evaluation rules from Sect. 3.3.

⁴ Non-deterministic constructs lead to semantical inconsistencies as one of the authors argues in [25]. They are not further discussed here.

4.1 Evaluation of Operation Contracts

The evaluation of an operation contract is defined with respect to a transition between two states.

StateTransition metaclass from our metamodel (see Fig. 3) is used to capture one transition from a pre- to a post-state. This transition is characteristic of one concrete operation execution with concrete values passed as operation parameters. In order to be able to evaluate one pre- or one post-condition we need all information about the state transition for which we want to perform the evaluation: operation that caused the transition, values of the operation parameters, pre-state, post-state, relationships between objects from pre- and post-state.

The evaluation of preconditions can be done analogously to the evaluation of invariants. The current state the evaluation rules referred to in the *Context* is in this case just the pre-state. In addition, the bindings for the operation arguments have to be extracted from a *Stimulus* that belongs to the *StateTransition* for which we perform the evaluation.

The evaluation of the postcondition is basically done in the post-state. The keyword *result* is evaluated according to the binding for the return parameter. The evaluation of *result* is fully analogous to the evaluation of variable expressions.

The evaluation of *@pre* is more complicated. It requires a switch between pre- and post-state, more precisely, we have to manage the different values for properties of each object in the pre- and post-state. Even more complicated, it might be the case that the set of objects itself has changed between pre- and post-state.

In the semantics of OCL described in [14, Annex A], the pre- and post-states are encoded as a set of functions (each function represents an attribute or a navigable association end) that work on a constant domain of objects. Furthermore, there is an extra function that keeps track which of the objects are created in the current state. The formalization has the advantage that the involved objects do not change their identity and thus is very easy to understand. Unfortunately, we were not able to apply this simple model to our semantics due to technical problems caused by the format of graph transformations. In our semantics, the objects in the pre- and post-state have different identities, but each object can be connected with one object from the opposite state via an instance of the *ObjectMap* metaclass. Please note, that for one object there can exist many *ObjectMaps* depending on number of *StateTransitions* one object is involved in. A pair of related objects represents the same object when we would view a pre-/post-state pair as an evolvment over the same domain. If an object from the pre-state is not related with any object from the post-state, it means that this object was deleted during the state transition. Analogously, objects in the post-state without a counterpart in the pre-state were created.

Fig. 25 shows an example. The pre-state consists of two objects with identifiers p1, p2 whose type is a class with name **Person**. The attribute

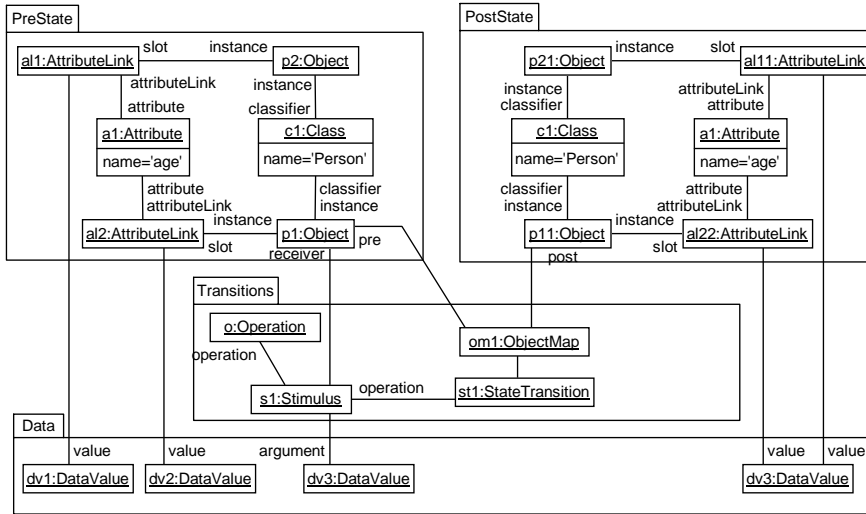


Fig. 25 Relationship between Pre- and Poststate

links for the attribute named `age` refer to the value `dv1` and `dv2`, which reside in the package `Data`. In the post-state, the identifiers for objects and attribute links have completely changed. But since object `p1` and `p11` are related by an *ObjectMap* `om1`, we know that `p11` and `p1` represent the same object. Note, however, that the state of this object has changed since the attribute link for attribute named `age` doesn't refer any longer to the value `dv2` but to `dv3`. Since there are no other *ObjectMaps* we can conclude that during the state transition from the pre-state to the post-state, the object `p2` was deleted and object `p21` was created.

The *@pre*-Operator can now be realized as an extension to the already existing core rules. Note that the official OCL syntax allows to attach *@pre* on every functor, but *@pre* is only meaningful when attached to Navigation Expressions or to an *allInstances*-expression. The most complicated case is the application to *AssociationEndCallExps*.

Figure 26 shows the extended evaluation rule for *AssociationEndCallExp* with a object-valued multiplicity (upper limit is 1). The current OCL metamodel encodes *@pre* expressions as operation call expressions of a pre-defined operation with name *@pre*. The source expression of this operation call expression is exactly that expression, to which the *@pre* operator is attached. The rule reads as follows: First, we wait for the situation in which the source expression of the association end call expression is evaluated (here, to `o1`). Note that the *Context* requires that `o1` is an object from the post-state (what should be always the case). Then, the corresponding object of `o1` in the pre-state is searched (`o1pre`) for which the original rule for evaluation of the association end call is applied (in the pre-state). The object representing the result of the association end call (`o2pre`) is then projected to the post-state (`o2`), what is then given back as the result of

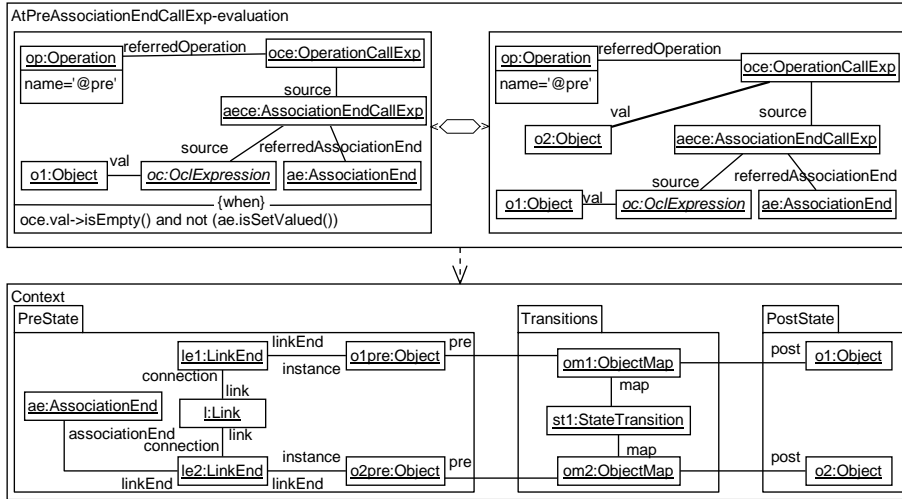


Fig. 26 Evaluation of @pre attached to an object-valued association end call expression

the evaluation. Note that we didn't specify so far the cases, in which `o1` does not have a counterpart on the pre-state (i.e. the source expression `oc` evaluates to a newly created object) or that the result of the association end call in the pre-state (`o2pre`) does not have a counterpart in the post-state (i.e. the object `o2pre` was deleted during the state transition). This question is answered in the next subsection.

4.2 Evaluation to undefined

The evaluation of OCL expressions to *undefined* is probably one of the most complicated semantic concepts in OCL and has raised many discussions. The value *undefined* has been often mixed in the literature with the null-value (known from Java). Furthermore, questions like *Can an AttributeLink refer to undefined in a state?* *Can a Set-expression be evaluated to undefined?* *Can a Set-value have elements that are undefined?* are not fully clarified by the official OCL semantics (cmp. also [24]).

First of all, we should note that the value *undefined* was added to the semantic domain for the sole purpose to indicate exceptional situations during the evaluation. For instance, when an object-valued *AssociationEndCallExp* tries to navigate over non-existing links or that a cast of an expression to a subclass fails. Thanks to the pre-defined operation `oclIsUndefined()` it is possible to test if an expression is currently evaluated to *undefined*; what – together with the exception from strict evaluation for *and*, *or*, *implies*, *forAll* etc – is a powerful tool to write OCL constraints reflecting the intended semantics even in the presence of *undefined* values.

But when is actually an expression evaluated to *undefined*? Strictly speaking, we had to add for each core evaluation rule a variant of this rule, that captures all situations in which undefinedness would occur. Fortunately, we have designed our evaluation rule in such a way, that this additional rule can be generated. Evaluation to *undefined* is always needed in all cases, in which the pattern given in the Context does not match with the current situation.

Let's have a look to the rule for *@pre* on association end call expressions (Fig. 26). If for instance of the object *o1* (evaluation of the source expression) was newly created during the state transition so that the pre-post link to an object *o1pre* is missing, then the whole *@pre*-Expression evaluates to undefined. Likewise, if the corresponding object *o1pre* exists but does not have a link for association end *ae*. Another reason could be that the link exist but the referred object *o2pre* is deleted during the state change. In all these cases, the *@pre*-Expression should be evaluated to *undefined* and these cases have in common that the pattern given in the *Context* does not match.

4.3 Dynamic Binding

Dynamic (or late) binding is one of the key concepts in object-oriented programming languages but has been mostly ignored in the OCL literature. Dynamic binding becomes relevant for the evaluation of user-defined queries. Let's assume we have two classes *A* and *B*, the class *B* is a subclass of *A* and the operation *m()* is declared as query with return type *Integer* in *A*.

We have the following constraints:

```
context A::m(): Integer
body: 5
```

```
context B::m(): Integer
body: 7
```

Let *a* and *b* be expressions that evaluate to an *A* and a *B* object, respectively. The result of the evaluation of *a.m()* is clearly 5. The evaluation of *b.m()* depends on whether or not OCL supports dynamic binding.

The core rule for query evaluation shown in Fig. 15 does not realize dynamic binding so far because it doesn't take into account potential inheritance hierarchy in the model. Result of the second rule shown in the figure is value of any body expression (*ocb2*) regardless its context.

For the situation when different bodies can be attached to the same operation (as in our example with classes *A* and *B*) we have to define a strategy for choosing the right body. The most suitable strategy would be to search the inheritance tree and take the body expression defined for the classifier that is the least parent of the source classifier (in the case of *b.m()* that would be the second body constraint 7).

In order to transform the static-binding evaluation rules for queries shown in Fig. 15 to a dynamic-binding rule, we had to alter the when-clauses in the LHS of the second rule with the following constraint:

```

if bm.booleanValue <> true then i.oclIsTypeOf(OclVoidValue)
else i=op.getRightBody(opce.source.val.oclAsType(Object)
    .classifier ->any(true))
endif and
c.stereotype.name='pre'

```

The *getRightBody* query (when multiple inheritance is not allowed) is defined as:

```

context Operation def: getRightBody(cl:Classifier):Instance
=
if self.body.oclAsType(ExpressionInOcl).contextualClassifier
    ->exists(cl) then
op.body->select(b|b.oclAsType(ExpressionInOcl)
    .contextualClassifier ->includes(cl))
    ->any(true).bodyExpression.val
else if cl.getDirectParent()->notEmpty() then
    self.getRightBody(cl.getDirectParent()->any(true))
    else getOclVoidValue()
    endif
endif

```

5 Tailoring OCL for DSLs

This section contains an example how our approach for defining the semantics of OCL can be applied for the definition of an OCL-based constraint language that is tailored to a domain specific language (DSL).

As a running example we will use a simple Relational Database Language for which we will define an extension of OCL. Two tables **Person** and **Dog** (see Fig. 27) will be used as an example, for which we develop domain-specific constraints. Each table has one primary key (**personID** for the **Person** table and **dogID** for the **Dog** table). In addition, column **ownerID** of table **Dog** has a foreign key relationship with the **personID** column of the **Person** table.

A simple metamodel for relational databases is shown in Fig. 28. This language is sufficient to specify the database from Fig. 27. Please note that, for the sake of simplicity, we have avoided to introduce database-specific types, but reuse already existing UML/MOF primitive types as types for table columns.

When tailoring OCL as a constraint/query language for a domain specific language, it is necessary to introduce additional concepts to OCL in order to capture domain specific constructs. In our example, two constructs require an extension of the OCL metamodel: 1) navigation to a column 2) navigation to a column constrained with a foreign key. The first navigation

| Person | | | Dog | | |
|------------------|-------|-----|---------------|----------|------------------------------|
| personID (PK) | name | age | dogID (PK) | breed | ownerID (FK for personID) |
| 1 | John | 23 | 1 | Doberman | 1 |
| 2 | Mark | 17 | 2 | Bulldog | 1 |
| 3 | Steve | 45 | 3 | Poodle | 2 |

Fig. 27 An Example of Relational Database

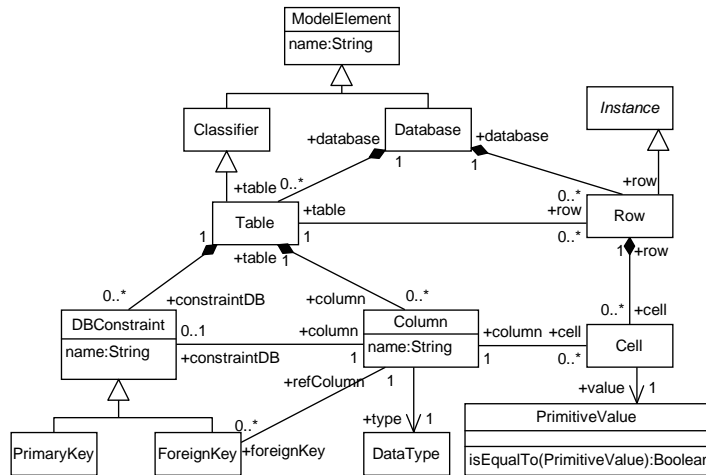


Fig. 28 Relational Database Metamodel

is applied on a *Row* and has to return the value of the *Column* for this *Row* and the second one has to return a *Row* of the *Table* to which the *ForeignKey* refers.

An example for these two new navigation expressions is the following:

```
Dog.allInstances()->select(d|d.breed='Doberman')
->forAll(dd|dd<=>ownerID.age>18)
```

This example expression uses three specificities of our relational database DSL: Ordinary navigation to columns *breed* and *age*, foreign key navigation to column *ownerID* (foreign key navigation is marked with $\langle = \rangle$ in order to make it distinguishable from ordinary column navigation), and a call of *allInstances()* on a table.

Another way of expressing the same could be by using only ordinary column navigation and *allInstances()*, but this version is much longer:

```
Dog.allInstances()->select(d|d.breed='Doberman')
->forAll(dd|Person.allInstances()
->any(p|p.personID=dd.ownerID).age>18)
```

In order to incorporate ordinary and foreign key column navigation into the constraint language, the metamodel for OCL had to be altered. Figure 29 shows the part of the Domain Specific Query language that is different from the standard OCL.

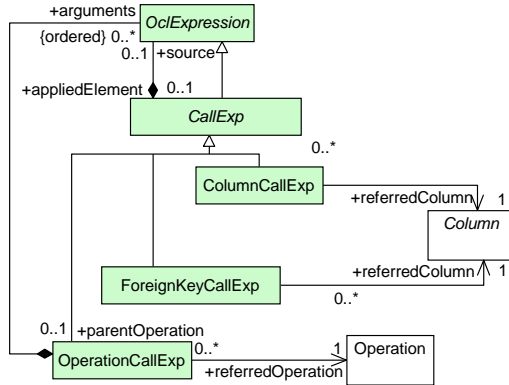


Fig. 29 DSL Navigation Expressions

Fig. 30 shows the definition of the semantics of column call expressions in form of an evaluation rule. The result of evaluation of such an expression would be value of the *Cell* that belongs to the *Row* that is the source of the expression, and that is referred by the chosen *Column*.

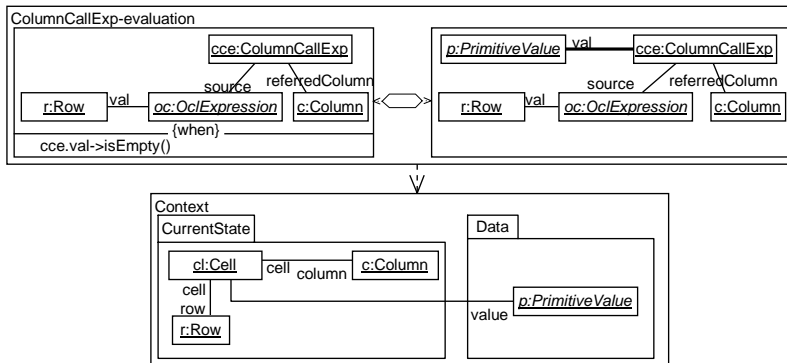


Fig. 30 Semantics of Column Navigation Specified with QVT

The semantics of *ForeignKeyCallExp* is shown in Fig. 31. This rule specifies that the value of the *ForeignKeyCallExp* will be a *Row* *r2* for which its primary key column has a *Cell* with the same value as the *Cell* of the source *Row* *r* for the foreign key column.

A mandatory construct that is needed when specifying the semantics of domain specific query languages and that cannot be reused from stan-




Fig. 31 Semantics of Foreign Key Navigation Specified with QVT

standard OCL is the operation call expression for the predefined operation `allInstances()`. This construct operates on model elements that do not exist in UML/MOF and therefore has to be explicitly defined as in Fig. 32.

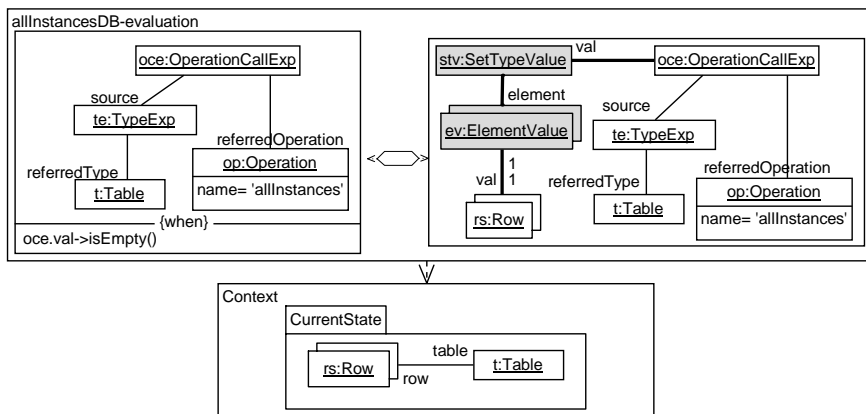


Fig. 32 Semantics of `allInstances` Operation Call Expression for Relational Database

Another way of defining the semantics of OCL expressions on the instance level is by moving (transforming) an OCL expression to an equivalent expression that queries the corresponding metamodel. As an example, consider the following `ColumnCallExpression` specified using our concrete syntax:

`exp.age`

Please note that the source expression `exp` can be any expression of type *Table*. This short expression in the DSL-specific version of OCL can be emulated by the following expression, which exploits the metalevel. However, this expression is clearly much more complicated.

```
Column.allInstances()->select(col|col.name=='age' and
col.table=exp.table).cell
    ->select(cc|cc.row=exp)
    ->any(true)
```

6 Related Work

The work described in this paper combines techniques and results from different fields in computer science: logics, precise modeling with UML/OCL, model transformation, modeling language design. For this reason, we separated related work into three categories.

6.1 Approaches to define the semantics of OCL

There are numerous papers and even some dissertations that propose a formal semantics for complete OCL or for a fragment of it, e.g., [26,27,15,28,22,29–32] and, recently, [21]. Many other papers have identified inconsistencies in the official OCL semantics and contributed in this form to a better understanding of OCL's concepts, e.g., [33,34,25,35,24].

Though we hope to have addressed in our semantics many of the issues raised in previous paper, there is no guarantee we can give, that our semantics has resolved all problems (a discussion on this would deserve another paper). What is more relevant for the current paper is to compare *the technique*, which has been used for the semantics definition, with that of other approaches. We restrict ourselves to a comparison with the two semantics given in the OCL language standard.

6.1.1 Official OCL Semantics: Informative Annex A of [14] presents a set-theoretical semantics for OCL, which goes back of the dissertation of Mark Richters [15]. This semantics has been marked in the OCL standard as *informative*.

The *semantic domain* of OCL is formalized by the notion of *system state* (a triple consisting of the set of objects, the set of attribute values for the objects, and the set association links connecting objects) and the *interpretation of basic types*. The notion of *system state* is defined on top of the notion of *object model*. What was formalized by Richters as *system state* is known in UML terminology as *object diagram*, an *object model* corresponds to a *class diagram*.

In our approach, the class and object diagrams are directly formalized by their metamodels and the interpretation of basic types is covered by the package *Values* of the OCL metamodel. All three metamodels, on which our approach relies, are part of the official language definition for UML/OCL. However, there is one important difference to Richter’s semantics: In Richter’s approach, one object can be in multiple states, whereas in our approach, states are represented by object diagrams which can never contain objects with the same identity. We solved this problem by introducing `ObjectMap` objects (cmp. Sect. 2.2) whenever two different states are involved in the evaluation of OCL constraints (e.g., post-conditions). Note that a set of `ObjectMap` objects referring to a pre-state and a post-state can also encode the information which of the objects were created/deleted during the transition from pre- to post-state. In Richter’s approach, the lifetime of an object is encoded by the function σ_{CLASS} .

The evaluation of OCL expressions is formalized in Richter’s semantics by an *interpretation function* \mathcal{I} , which is defined separately for each type of OCL expression. The definitions for \mathcal{I} are based on the above mentioned ingredients of the semantics *object model*, *system state*, *interpretation of basic types*. In our approach, the interpretation function \mathcal{I} is implicitly given by QVT rules, which are based on the metamodels for class diagrams, object diagrams, and on the *Values* package.

One of the most interesting details when comparing the formalization of expression evaluation is the handling of pre-defined functions. Following Richter, pre-defined functions like `=`, `union`, `concat`, etc., are interpreted by their mathematical counterparts, e.g. $\mathcal{I}(=_t)(v_1, v_2) = true$ if $v_1 = v_2$ and $v_1 \neq \perp$ and $v_2 \neq \perp$. Otherwise stated, the semantics of some operations of the object language (OCL) is reduced to the semantics of some operations of the meta language (mathematics). The same holds in our case, the semantics of operation `'='` of the object language (OCL) is reduced to the semantics of the operation `'='` in the metalanguage (QVT) (see Sect.3.3.2).

In both cases, it has to be assumed that the semantics of the meta-language has been already defined *externally* (cmp. also [36]). In case of Richter’s semantics, one could refer to textbooks introducing mathematics. In case of our semantics, we can refer to the implementation of QVT engines, which actually map QVT rules to statements in a programming language, e.g. Java.

6.1.2 Official OCL Semantics: Normative The semantics described in [14], Sect. 10 *Semantics Described Using UML* is called *normative OCL semantics* and shares the same main goal as our approach: to have a semantics description of OCL, which is seamlessly integrated into the other artifacts (metamodels) of OCL’s language definition. However, there are important differences.

The normative semantics defines a package *Values* to encode pre-defined data types and system states. We tried to align our approach as much as possible with this *Values* package (e.g. `NameValueBinding`), but some details

differ. Most notable, as already mentioned in the comparison with Richter’s semantics, our states never contain identical objects. The normative OCL semantics insists on keep object identities across states, but this yields to a quite complicated encoding of attribute value and links, which have to be kept separated from objects (see metaclass `LocalSnapshot`). Moreover, the normative semantics encodes exactly one system trace (metaassociation `pred--succ` on `LocalSnapshot`), while in our approach state transitions are modeled explicitly by a new metaclass `StateTransition`.

The evaluation of OCL expressions is formalized in the normative semantics by so-called *evaluation classes*. For each metaclass from the metamodel of OCL’s abstract syntax, there is exactly one corresponding evaluation class, e.g. `AttributeCallExpEval`. Evaluation classes are complimented by a number of invariants, whose purpose is to specify the evaluation process. In many cases, the invariants can be mapped to exactly one QVT rule in our approach. For example, there is for each evaluation class one invariant specifying the propagation of the current binding of variables (called `Environment` in the normative semantics) to sub-expressions, what corresponds to our variable binding propagation rules described in Sect. 2.4.

The normative semantics has been also the starting point for a semantics formalization given by Chiaradía and Pons in [37]. They alter the OCL semantics’ metamodel by introducing visitor pattern in order to reduce the duplication of information in *AbstractSyntax* and *Evaluations* packages of OCL metamodel. Contrary to our approach, they use UML sequence diagrams to express the semantics of OCL expressions.

6.2 Approaches to define language semantics by model transformations

The application of model transformations (or, more general, graph transformations) for the purpose of defining language semantics is not a new idea. However, we are only aware of one paper, which applies this technique for the definition of the semantics of OCL. Bottoni et al. propose in [38] a graphical notation of OCL constraints and, on top of this notation, some simplification rules for OCL constraints. These simplification rules specify implicitly the evaluation process of OCL expressions. However, the semantics of OCL is not developed as systematically as in our approach, only the simplification rules for *select* are shown. Since [38] was published at a time where OCL did not have an official metamodel, the simplification rules had to be based on another language definition of OCL.

For behaviorial languages, Engels et al. define in [39] a dynamic semantics in form of graph-transformation rules, which are similar to our QVT rules. As an example, the semantics of UML statechart diagrams is presented.

In [40] Varró points out the abstraction gap between the ”graphical” world of UML and mathematical models used to describe dynamic semantics. In order to fill this gap he uses graph transformation systems to describe

visual operational semantics. Application of the approach is demonstrated by specifying semantics of UML statecharts.

Stärk et al. define in [41] a formal operational semantics for Java by rules of an Abstract State Machine (ASM). The semantic domain of Java programs is fixed by defining the static structure of an appropriate ASM. The ASM encodes furthermore the Abstract Syntax Tree (AST) of Java programs. As shown by our motivating example in Sect. 2, there are no principal differences between an AST and an instance of the metamodel. Also, ASM and QVT rules are based on the same mechanisms (pattern matching and rewriting).

6.3 Other related work

An interesting classification of OCL language concepts was developed by Chiorean et al. in [42]. In this paper, OCL language constructs are classified according to their usage in different domains, such as *Transformations*, *Assertions*, and *Commands*. In our approach, we have concentrated on what is called *core OCL* in [42], but it would be definitely worthwhile to investigate the other domains as well.

Kolovos et al. define in [43] a navigation language for relational databases that is similar to our language defined in Sect. 5. They use the metalanguage *EOL* (which is based on OCL) to define the result of evaluation of new expressions like column navigation.

7 Conclusions and Future Work

We have developed a metamodel-based, graphical definition of the semantics of OCL. Our semantics consists of a metamodel of the semantic domain (we have slightly adapted the existing metamodels from UML1.x) and a set of transformation rules written in an extension of QVT that specify formally the evaluation of an OCL constraint in a snapshot. To read our semantics, one does not need advanced skills in mathematics or even knowledge in formal logic; it is sufficient to have a basic understanding of metamodeling and QVT. The most important advantage, however, is the flexibility our approach offers to adapt the semantics of OCL to domain-specific needs. Since the evaluation rules can directly be executed by any QVT compliant tool, it is now very easy to provide tool support for a new dialect of OCL. This is an important step forward to the OMG's vision to treat OCL as a *family of languages*.

We are currently investigating how an OCL semantics given in form of QVT rules can be used to argue on the semantical correctness of refactoring rules for UML/OCL, which we have defined as well in form of QVT rules. A refactoring rule describes small changes on UML class diagrams with attached OCL constraints. A rule is considered to be *syntactically correct* if in all applicable situations the refactored UML/OCL model is syntactically

well-formed. We call a rule *semantically correct* if in any given snapshot the evaluation of the original OCL constraint and the refactored OCL constraint yields to the same result (in fact, this view is a simplified one since the snapshots are sometimes refactored as well). To argue on semantical correctness of refactoring rules, it has been very handy to have the OCL semantics specified in the same formalism as refactoring rules, in QVT. A more detailed description together with a complete argumentation on the semantical correctness of the *MoveAttribute* refactoring rule can be found in [44].

Another branch of future activities is the description of the semantics of programming languages with graphical QVT rules. Our ultimate goal is to demonstrate that also the description of the semantics of a programming language can be given in an easily understandable, intuitive format. This might finally contribute to a new style of language definitions where the semantics of the language can be formally defined as easy and straightforward as it is today already the case with the syntax of languages.

References

1. OMG. UML 2.0 Infrastructure Specification. OMG Document ptc/03-09-15, Sep 2003.
2. Kirsten Berkenkötter. OCL-based validation of a railway domain profile. In Thomas Kühne, editor, *Models in Software Engineering, Workshops and Symposia at MoDELS 2006, Genoa, Italy, October 1-6, 2006, Reports and Revised Selected Papers*, volume 4364 of *Lecture Notes in Computer Science*, pages 159–168. Springer, 2007.
3. Birgit Demuth, Heinrich Hußmann, and Sten Loecher. OCL as a specification language for business rules in database applications. In *UML'01: Proceedings of the 4th International Conference on The Unified Modeling Language, Modeling Languages, Concepts, and Tools*, volume 2185 of *Lecture Notes in Computer Science*, pages 104–117. Springer, 2001.
4. Eric Cariou, Raphaël Marvie, Lionel Seinturier, and Laurence Duchien. OCL for the specification of model transformation contracts. In Octavian Patrascioiu, editor, *OCL and Model Driven Engineering, UML 2004 Conference Workshop, October 12, 2004, Lisbon, Portugal*, pages 69–83. University of Kent, 2004.
5. David H. Akehurst and Behzad Bordbar. On querying UML data models with OCL. In Martin Gogolla and Cris Kobryn, editors, *UML 2001 - The Unified Modeling Language, Modeling Languages, Concepts, and Tools, 4th International Conference, Toronto, Canada, October 1-5, 2001, Proceedings*, volume 2185 of *Lecture Notes in Computer Science*, pages 91–103. Springer, 2001.
6. Birgit Demuth and Heinrich Hußmann. Using UML/OCL constraints for relational database design. In Robert B. France and Bernhard Rumpe, editors, *UML'99: The Unified Modeling Language - Beyond the Standard, Second International Conference, Fort Collins, CO, USA, October 28-30, 1999, Proceedings*, volume 1723 of *Lecture Notes in Computer Science*. Springer, 1999.

7. Hanna Bauerdick, Martin Gogolla, and Fabian Gutsche. Detecting OCL traps in the UML 2.0 superstructure: An experience report. In Thomas Baar, Alfred Strohmeier, Ana M. D. Moreira, and Stephen J. Mellor, editors, *UML 2004 - The Unified Modelling Language: Modelling Languages and Applications. 7th International Conference, Lisbon, Portugal, October 11-15, 2004. Proceedings*, volume 3273 of *Lecture Notes in Computer Science*, pages 188–196. Springer, 2004.
8. Oslo Team. Oslo project. <http://oslo-project.berlios.de/>, 2007.
9. MDT-OCL Team. Eclipse MDT - OCL project. <http://www.eclipse.org/modeling/mdt/?project=ocl>, 2007.
10. Dresden OCL Team. Dresden OCL Toolkit. <http://dresden-ocl.sourceforge.net/>, 2007.
11. Octopus Team. Octopus - OCL Tool for Precise Uml Specifications. <http://octopus.sourceforge.net/>, 2007.
12. USE Team. USE - a UML-based Specification Environment. <http://www.db.informatik.uni-bremen.de/projects/USE/>, 2007.
13. OCLE Team. OCLE - Object Constraint Language Environment. <http://lci.cs.ubbcluj.ro/ocle/index.htm>, 2007.
14. OMG. Object Constraint Language - OMG Available Specification, version 2.0. OMG Document formal/06-05-01, May 2006.
15. Mark Richters. *A precise approach to validating UML models and OCL constraints*. PhD thesis, Bremer Institut für Sichere Systeme, Universität Bremen, Logos-Verlag, Berlin, 2001.
16. OMG. Meta object facility (MOF) 2.0 Query/View/Transformation Specification. OMG Document ptc/05-11-01, Nov 2005.
17. Slaviša Marković. Evaluating UML/OCL using QVT. <http://lgl.epfl.ch/members/markovic/evaluation/>, 2007.
18. RocLET Team. RocLET project. <http://www.roclet.org/>, 2007.
19. Borland. Together technologies. www.borland.com/together/, 2007.
20. Achim D. Brucker and Burkhard Wolff. The HOL-OCL book. Technical Report 525, ETH Zurich, 2006.
21. Achim D. Brucker. *An Interactive Proof Environment for Object-oriented Specifications*. Ph.D. thesis, ETH Zurich, 2007. ETH Dissertation No. 17097.
22. Tony Clark, Andy Evans, and Stuart Kent. Engineering modelling languages: A precise meta-modelling approach. In Ralf-Detlef Kutsche and Herbert Weber, editors, *Fundamental Approaches to Software Engineering. 5th International Conference, FASE 2002 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2002 Grenoble, France, April 2002, Proceedings*, volume 2306 of *LNCS*, pages 159–173. Springer, 2002.
23. Slaviša Marković and Thomas Baar. An OCL semantics specified with QVT. In Oscar Nierstrasz, Jon Whittle, David Harel, and Gianna Reggio, editors, *Proceedings, MoDELS/UML 2006, Genova, Italy, October 1-6, 2006*, volume 4199 of *LNCS*, pages 660–674. Springer, October 2006.
24. Achim D. Brucker, Jürgen Doser, and Burkhard Wolff. Semantic issues of OCL: Past, present, and future. In Birgith Demuth, Dan Chiorean, Martin Gogolla, and Jos Warmer, editors, *OCL for (Meta-)Models in Multiple Application Domains*, pages 213–228, Dresden, 2006. University Dresden. Available as Technical Report, University Dresden, number TUD-FI06-04-Sept. 2006.
25. Thomas Baar. Non-deterministic constructs in OCL - what does any() mean. In Andreas Prinz, Rick Reed, and Jeanne Reed, editors, *Proc. 12th SDL Forum, Grimstad, Norway, June 2005*, volume 3530 of *LNCS*, pages 32–46. Springer, 2005.

26. Mark Richters and Martin Gogolla. On formalizing the UML object constraint language OCL. In Tok Wang Ling, Sudha Ram, and Mong Li Lee, editors, *Proc. 17th Int. Conf. Conceptual Modeling (ER'98)*, pages 449–464. Springer, Berlin, LNCS 1507, 1998.
27. Mark Richters and Martin Gogolla. A metamodel for OCL. In Robert France and Bernhard Rumpe, editors, *UML'99 - The Unified Modeling Language. Beyond the Standard. Second International Conference, Fort Collins, CO, USA, October 28-30. 1999, Proceedings*, volume 1723 of *LNCS*, pages 156–171. Springer, 1999.
28. María Victoria Cengarle and Alexander Knapp. A formal semantics for OCL 1.4. In Martin Gogolla and Cris Kobryn, editors, *UML*, volume 2185 of *Lecture Notes in Computer Science*, pages 118–133. Springer, 2001.
29. Stephan Flake and Wolfgang Mueller. Formal Semantics of Static and Temporal State-Oriented OCL-Constraints. *Journal on Software and System Modeling (SoSym)*, 2(3):164.186, 2003.
30. Rolf Hennicker, Alexander Knapp, and Hubert Baumeister. Semantics of OCL operation specifications. *Electronic Notes in Theoretical Computer Science, Proceedings of OCL 2.0 Workshop at UML'03*, 102:111–132, 2004.
31. Thomas Baar. *Über die Semantikbeschreibung OCL-artiger Sprachen*. PhD thesis, Fakultät für Informatik, Universität Karlsruhe, 2003. ISBN 3-8325-0433-8, Logos Verlag, Berlin, In German.
32. María Victoria Cengarle and Alexander Knapp. OCL 1.4/5 vs. 2.0 expressions formal semantics and expressiveness. *Software and System Modeling*, 3(1):9–30, 2004.
33. Steve Cook, Anneke Kleppe, Richard Mitchell, Bernhard Rumpe, Jos Warmer, and Alan Cameron Wills. The amsterdam manifesto on OCL. In Tony Clark and Jos Warmer, editors, *Object Modeling with the OCL: The Rationale behind the Object Constraint Language*, pages 115–149. Springer, 2002.
34. Stephan Flake. Ocltype - a type or metatype? *Electr. Notes Theor. Comput. Sci.*, 102:63–75, 2004.
35. David H. Akehurst, Gareth Howells, and Klaus D. McDonald-Maier. Supporting OCL as part of a family of languages. In Thomas Baar, editor, *Proceedings of the MoDELS'05 Conference Workshop on Tool Support for OCL and Related Formalisms - Needs and Trends, Montego Bay, Jamaica, October 4, 2005*, Technical Report LGL-REPORT-2005-001, pages 30–37. EPFL, 2005.
36. David Harel and Bernhard Rumpe. Meaningful Modeling: What's the Semantics of "Semantics"? *IEEE Computer Software*, 37(10):64–72, October 2004.
37. Juan Martín Chiaradía and Claudia Pons. Improving the OCL semantics definition by applying dynamic meta modeling and design patterns. In Birgith Demuth, Dan Chiorean, Martin Gogolla, and Jos Warmer, editors, *OCL for (Meta-)Models in Multiple Application Domains*, pages 229–239, Dresden, 2006. University Dresden. Available as Technical Report, University Dresden, number TUD-FI06-04-Sept. 2006.
38. Paolo Bottoni, Manuel Koch, Francesco Parisi-Presicce, and Gabriele Taentzer. Consistency checking and visualization of OCL constraints. In Andy Evans, Stuart Kent, and Bran Selic, editors, *UML 2000 - The Unified Modeling Language, Advancing the Standard, Third International Conference, York, UK, October 2-6, 2000, Proceedings*, volume 1939 of *LNCS*, pages 294–308. Springer, 2000.

39. Gregor Engels, Jan Hendrik Hausmann, Reiko Heckel, and Stefan Sauer. Dynamic meta modeling: A graphical approach to the operational semantics of behavioral diagrams in UML. In Andy Evans, Stuart Kent, and Bran Selic, editors, *UML 2000 - The Unified Modeling Language, Advancing the Standard, Third International Conference, York, UK, October 2-6, 2000, Proceedings*, volume 1939 of *LNCS*, pages 323–337. Springer, 2000.
40. Dániel Varró. A formal semantics of UML Statecharts by model transition systems. In Andrea Corradini, Hartmut Ehrig, Hans-Jörg Kreowski, and Grzegorz Rozenberg, editors, *Proc. ICGT 2002: 1st International Conference on Graph Transformation*, volume 2505 of *LNCS*, pages 378–392, Barcelona, Spain, October 7–12 2002. Springer-Verlag.
41. Robert F. Stärk, Joachim Schmid, and Egon Börger. *Java and the Java Virtual Machine - Definition, Verification, Validation*. Springer, 2001.
42. Dan Chiorean, Maria Bortes, and Dyan Corutiu. Proposals for a widespread use of OCL. In Thomas Baar, editor, *Tool Support for OCL and Related Formalisms - Needs and Trends, MoDELS'05 Conference Workshop, Montego Bay, Jamaica, October 4, 2005, Proceedings*, Technical Report LGL-REPORT-2005-001, pages 68–82. EPFL, 2005.
43. Dimitrios S. Kolovos, Richard F. Paige, and Fiona A.C. Polack. Towards using OCL for instance-level queries in domain specific languages. In Birgith Demuth, Dan Chiorean, Martin Gogolla, and Jos Warmer, editors, *OCL for (Meta-)Models in Multiple Application Domains*, pages 26–37, Dresden, 2006. University Dresden. Available as Technical Report, University Dresden, number TUD-FI06-04-Sept. 2006.
44. Thomas Baar and Slaviša Marković. A graphical approach to prove the semantic preservation of UML/OCL refactoring rules. In Irina Virbitskaite and Andrei Voronkov, editors, *Proceedings, Sixth International Andrei Ershov Memorial Conference on Perspectives of System Informatics (PSI 2006), Akademgorodok near Novosibirsk, Russia*, volume 4378 of *LNCS*, pages 70–83. Springer, 2007.