

Certification of Transformation Algorithms in Model-Driven Software Development

Miguel Garcia, Ralf Möller

Institute for Software Systems (STS)
Hamburg University of Technology (TUHH), D-21073 Hamburg
{miguel.garcia, r.f.moeller}@tuhh.de

Abstract: The increasing reliance on Model-Driven Software Development calls for model compilers to assume the role of today's compilers, i.e., reliability of these components is of utmost importance. We describe how to certify model transformations in this context by bridging the gap between the languages in which such transformations are specified (e.g., Essential MOF, OCL, OO programs) and the decision procedures needed to verify properties expected of such transformations. Two major aspects are investigated in this paper: (i) valid output is obtained for each valid input, (ii) the output satisfies certain properties. Results from application projects validate our approach, which internally applies model-driven techniques to the certification process itself by mapping transformation specifications into the ⁺CAL model-checking language.

1 Introduction

Model-Driven Software Development (MDS) is gaining consensus in the Software Engineering community as a viable technology to improve both productivity and quality. Metamodel-based definitions of Domain Specific Languages (DSLs) are routinely leveraged to automatically derive language processing tools (syntax-aware editors [JBK06], diagram editors [EEHT05], software repositories [APM03]) using techniques that build upon a common infrastructure such as the Eclipse Modeling Framework (EMF) [BSM⁺03]. The sketched productivity gains in jumpstarting a toolchain for MDS can be traced back to the specification of the static semantics (also called well-formedness rules, WFRs, or validity-checking rules) of DSLs in terms of declarative OCL invariants [WK03] over an Essential MOF (EMOF) metamodel [OMG06]. Such metamodels contain a wealth of machine-processable information, relieving tool implementors from manually hardwiring such specs into language processing tools (e.g., in the semantic analysis phase of a dedicated DSL compiler).

In order for MDS to deliver on its full potential, commensurate progress is required to increase the quality of emerging model transformers and compilers. A methodology and its associated tooling is presented in this paper to reach that goal for a representative class of model transformations. In particular, we investigate languages for which an EMOF + OCL metamodel is available. This methodology involves the automatic translation of the

input and output metamodels into a formalism for which a decision procedure is available to answer whether a given procedural transformation exhibits certain properties of interest. Two basic desirable guarantees for such transformations are (a) that all output sentences belong to the target language [HZS05], and (b) that the transformation function covers the whole input language for which it was designed [WKC06]. Experiences with current model-driven tooling shows that these basic requirements are not always met. Beyond these general requirements, guarantees specific to a given transformation are also desirable. For example, that an optimized implementation should produce the same result as the non-optimized version.

Nowadays, model compilers are in operation for various kinds of application tasks. For instance, transformations into Java Enterprise Edition (Java EE, defined in JSR-220) are very popular and include: DASL [Gol05], SecureUML [BDW06], and WebML [Cer05]. These DSLs allow for the specification of three-tier enterprise systems at a high level of abstraction and have metamodel-based descriptions. Similarly, platform-specific metamodels are available. For example, the persistent query language of Java EE (EJB3QL) has been metamodeled in [Gar06]: all the normative restrictions formulated in English in Ch. 4 of JSR-220 are recast as OCL invariants. Instantiations of the EJB3QL metamodel satisfying those invariants are valid abstract syntax trees (ASTs) for particular EJB3QL queries, and can be unparsed into text. Along the same lines, a complete metamodel of BPEL 1.1 is discussed in [Ake04].

In the current state of the practice, the WFRs on the input and output ASTs are evaluated at transformation-time, for each application of the transformation. We aim instead at certifying transformation algorithms at design-time, to make runtime-checks redundant, but moreover to get model compilers right early on, instead of patching them as new cases are discovered which were overlooked before widespread deployment.

As decision procedure for the task described above we adopt model-checking as described in [Lam06b]. A model-checker manipulates execution traces, which can be conceptualized as trees of states (each state also called a system snapshot), with root states derived to cover all initial conditions. Candidate successor states are computed from all actions enabled in the current state.

Once a transformation algorithm has been specified, at algorithm-design time the model-checker can detect situations where the transformation does not terminate, or terminates without establishing the properties of interest. The coverage achieved by the admittedly finite analysis of a model-checker is much higher than testing because (a) states can be manipulated symbolically, and (b) several properties of interest depend on the shape of an object graph rather than on its size or concrete attribute values (the “small scope hypothesis”) [Jac06]. For example, the condition “two lines intersect” can be manipulated without considering concrete crossing points. Model-checkers can detect these situations, taking a single state as representative of all those exhibiting such shape.

In summary, the contribution of this paper is twofold. First, we eliminate the laborious task of preparing a model-checkable specification for nontrivial MDSD transformations by reusing the EMOF/OCL WFRs for the static semantics contained in the metamodels of the languages participating in the transformation. Second, we demonstrate how to support the

development of robust transformation algorithms by employing a model-checking engine. In particular, we investigate the language ${}^+CAL$ and use the corresponding model-checker TLC. The significance of the approach is demonstrated using an application scenario concerning a well-known graph transformation problem (Schorr-Waite).

The structure of this paper is as follows. Sec. 2 elaborates on static semantics and derives set-theoretic definitions for EMOF + OCL metamodels. Afterwards, it is exemplified how OCL pre- and postconditions can be mapped to ${}^+CAL$. This is followed (in Sec. 3) by a discussion of the steps to follow when applying the proposed model-checking-based certification methodology. Afterwards, the section focuses on the analysis of a sample transformation. Alternative and complementary approaches to model-based verification are discussed in Sec. 4. Knowledge is assumed from the reader about MDS, set-theory and logic. The software artifacts discussed in this paper can be downloaded from [Gar].

2 Formalization of Essential MOF + OCL for model-checking

${}^+CAL$ [Lam06b] is a specification language designed to replace pseudo-code for writing high-level descriptions of algorithms. A ${}^+CAL$ algorithm manipulates mathematical objects in a series of steps. The granularity of a step is chosen by the algorithm designer, ranging from a single built-in statement to a composite step involving several ${}^+CAL$ statements. A step exhibits transaction semantics: intermediate results are not visible to concurrently executing processes and system invariants are required to hold only at step boundaries.

${}^+CAL$ includes control-flow statements typical from block-structured programming (if-then-else, while-do, sequential composition), as well as constructs for expressing non-deterministic and concurrent algorithms. For the purposes of this work, we focus on the sequential subset of ${}^+CAL$. Mathematical expressions and logical formulae may appear in ${}^+CAL$ programs whenever a construct calls for a value (e.g., in the condition part of an if-then-else, in an assert statement). In fact, the properties an algorithm should exhibit are routinely expressed as mathematical assertions on the input and output data (in our case: ASTs), with the model-checker being able to compute them for finite system snapshots.

The proposed certification methodology comprises the following steps:

1. Automatic translation of the definitions contained in the participating EMOF + OCL metamodels into ${}^+CAL$ (these definitions allow the transformation algorithm to refer to well-formed ASTs).
2. Expressing a model transformation as a ${}^+CAL$ algorithm operating on ASTs.
3. If appropriate, annotating the transformation with assumptions about its input, beyond the constraints expressed in the metamodels of input ASTs. The invoker of the transformation is responsible for satisfying these assumptions.
4. Annotating the transformation with assurances about the system state (in particular about the output ASTs) after every successful run of the transformation on valid

input (i.e. on well-formed input satisfying the assumptions made in the previous item). The algorithm is responsible for satisfying these assurances.

Following Design-By-Contract [MMS98], the assumptions in (3) are called *transformation preconditions* and the assurances in (4) *transformation postconditions*. This terminology overlaps with that of Hoare logic, where a postcondition fully specifies the transformation function between the pre and post states (no procedural statements are necessary). The checks that Design-By-Contract performs at runtime can be carried out at transformation-design time thanks to model-checking. As to the language for pre- and postconditions, besides ⁺CAL some or all of them can be expressed in OCL, with our prototype taking care of their translation.

2.1 Translating EMOF into ⁺CAL

This translation of AST structures results in ⁺CAL procedures to allocate instances and manipulate links and attributes, constituting a certified building block that simplifies the expression of the transformation algorithm. Our encoding of EMOF + OCL into ⁺CAL leverages previous work on set-theoretic semantics of UML and OCL [BKS02]. In the context of UML, a thorough analysis of the logical consistency of the MOF 2.0 specification is also reported in [AS06]. Our work focuses on the more recent EMOF and uses ⁺CAL instead [Lam06b]. Due to space restrictions we cannot present the transformation of all EMOF language constructs, but as an example for the main ideas Figure 2 shows the invariant for the bidirectionally constraint stated in the EMOF model in Figure 1



Figure 1: Sample EMOF model

The translation of ASTs results in ⁺CAL procedures for the allocation of instances and the manipulate of links and attributes. The generated specification allows for checking in every execution trace whether the constraints specified in the EMOF model (about multiplicities, ordering, bidirectionality, etc.) and all OCL invariants are maintained, and this for arbitrary finite system snapshots (being automatically generated by the model-checker).

2.2 Translating OCL into ⁺CAL

OCL prescribes that, at the end of a transaction, each system snapshot should fulfill the specified invariants. Lacking transaction demarcation, it is generally agreed that invariants

for each instance b of B , those instances of A reachable over D from it must in turn have b among their instances reachable over D

$$\begin{aligned} \text{InvariantBidirectionality_AtoB} &\triangleq \\ &\vee \text{AsForBoverD} = \langle \rangle \\ &\vee \forall b \in \text{DOMAIN AsForBoverD} : \\ &\quad \text{LET AsReachable} \triangleq \text{AsForBoverD}[b] \text{ IN } \quad \text{a sequence} \\ &\quad \forall i \in \text{DOMAIN AsReachable} : \\ &\quad \quad \text{LET anA} \triangleq \text{AsReachable}[i] \text{ IN} \\ &\quad \quad b \in \text{BsForAoverD}[anA] \end{aligned}$$

counterpart of the above, this time for each a

$$\begin{aligned} \text{InvariantBidirectionality_BtoA} &\triangleq \\ &\vee \text{BsForAoverD} = \langle \rangle \\ &\vee \forall a \in \text{DOMAIN BsForAoverD} : \\ &\quad \text{LET BsReachable} \triangleq \text{BsForAoverD}[a] \text{ IN } \quad \text{a set} \\ &\quad \forall aB \in \text{BsReachable} : \\ &\quad \quad \text{ElemIsInSeq}(a, \text{AsForBoverD}[aB]) \end{aligned}$$

$$\begin{aligned} \text{InvariantsDirectionality} &\triangleq \wedge \text{InvariantBidirectionality_AtoB} \\ &\quad \wedge \text{InvariantBidirectionality_BtoA} \end{aligned}$$

Figure 2: Bidirectionality expressed in ${}^+\text{CAL}$

in an OO program should hold after object construction and after the execution of each public operation. This concept of transaction boundaries is directly supported by ${}^+\text{CAL}$ step granularity. The *efficient* evaluation of invariants can be challenging. Usually an operation involves a small number of updates which leave most invariants unaffected. Being ${}^+\text{CAL}$ based on a temporal logic, there is no shortage of expressivity to encode OCL invariants, with the model-checker evaluating them in the background as execution traces are considered. Model-checkers can take into account the data-flow dependencies of formulae so as to approach non-redundant yet complete evaluation. This shortens the time elapsed from submitting an algorithm till counterexamples are found, thus increasing the productivity of the transformation designer. As already mentioned at the beginning of this section, OCL pre- and postconditions are translated as assertions into ${}^+\text{CAL}$, i.e., they are no substitute for the specification of the input-output transformation, which must be given as imperative statements.

The conversion from OCL to ${}^+\text{CAL}$ is performed by visitors over ASTs of OCL expressions, similar to the work of [BKS02]. A metamodel of ${}^+\text{CAL}$ was prepared for this purpose. This AST-to-AST conversion in turn is thus amenable to be verified with the techniques described in this paper. The inspiration to bootstrap the verification of this transformation was offered by the algorithm to translate from ${}^+\text{CAL}$ to TLA^+ (Temporal Logic of Actions [Lam02], the logic underlying ${}^+\text{CAL}$), which is itself specified in TLA^+ .

3 Certification process

Certification of an algorithm is an iterative process. Whenever, at design time it can be shown that an algorithm is bound to fail at runtime (i.e., for some inputs does not terminate, breaks metamodel invariants, or does not fulfill its part of the contract in establishing the postconditions) the model-checker not only indicates failure but presents an execution trace leading to that situation (a counterexample). The algorithm designer may apply a combination of (i) reformulating the algorithm to handle the situation that caused the failure, (ii) strengthening the preconditions (making the algorithm applicable to a subset smaller than well-formed ASTs), or (iii) weakening the postconditions. The practical limit to postcondition weakening is that the output must still be well-formed, as demanded by metamodel invariants.

3.1 Specifying algorithms in EMOF

As a further means to increase certification productivity, a textual syntax for an object-oriented programming language (“Executable EMOF”, or xEMOF for short) could define statements to manipulate instances of metamodels, as a high-level notation to express model transformation algorithms. The rationale for this is the large number of transformations already expressed in terms of the Visitor design pattern [GVJH98]. In a green field scenario, a language such as xEMOF would also prove useful by reducing the conceptual distance between a transformation algorithm and its implementation, thus shortening the certification process.

A language such as xEMOF is not just a thin layer of syntactic sugar over ⁺CAL. Instead, the translation is non-trivial because (a) method-dispatch in an object-oriented language depends on runtime-types instead of only declared types, (b) method overloading similarly complicates method selection, and (c) the interplay between inheritance and object initialization has to be taken into account. A more realistic modeling of Java, including for examples exceptions (which introduce alternative return paths and require bookkeeping to correctly unfold the activation frames in the call-stack) would not add expressive power yet complicate the translator.

Besides those transformations expressed in terms of visitors, another large group of existing transformations relies on pattern-matching mechanisms followed by in-place transformations (ATL [JK06], QVT [OMG05], graph-grammars [AKRS06]). Provided that they manipulate EMOF-based ASTs, their execution engines can similarly be formulated in terms of ⁺CAL reusing the translation performed by our prototype.

3.2 Certifying a non-trivial in-place transformation: Schorr-Waite

As to the readability of ⁺CAL, Figure 3 shows the encoding of the Schorr-Waite graph-marking algorithm. It is complex enough to serve as a reference case in source code

verification [HM05, MN05], as it involves modifying in-place an AST-like pointer-rich data structure, yet intuitive enough that its operation can be explained succinctly. To our knowledge, this is the first account on model-checking Schorr-Waite. Although we manually created the example specification from Figure 3, the discussion in the previous subsection and also, e.g., the results of [BB06] indicate that manual efforts can be reduced if not eliminated in the near future.

The Schorr-Waite algorithm performs a depth-first traversal of a directed graph, starting from a specific node called the root. Given that memory is at a premium during garbage collection, Schorr-Waite offers a constant upper bound on memory usage by avoiding keeping a stack with the nodes in the current path. Instead, as new nodes are visited, the link that was followed is reversed in place. Upon going back along the current path, the algorithm reconstructs the original topology. Pointer reversal avoids the introduction of a stack.

The correctness guarantees we expect are: (a) termination, (b) that all nodes reachable from the root (and only those) are marked, and (c) that the algorithm leaves the topology unchanged. These guarantees are encoded as assertions: the results computed by the implementation are compared with those computed from mathematical definitions which are executable in $^+$ CAL. For example, the set of nodes marked by Schorr-Waite is compared to the transitive closure over the `pointsTo` (mathematical) relationship applied to the root (our download [Gar] contains the full source code of this paper).

We model-checked the $^+$ CAL algorithm in Figure 3 using TLC and investigated graphs with up to 10 nodes and found that acceptable runtimes can be achieved (see also [Lam06a] for the empirical behavior of TLC). Thus, we conclude that the proposed methodology is practically significant. As mentioned before, due to the small scope hypothesis [Jac06], many, if not all, problems will be found using this problem size.

4 Related work: Alternative and complementary approaches

The methodology reported in this paper allows for the validation of model transformation algorithms. If implemented carefully, the assertions made for an algorithm carry over to its implementation. This manual step introduces the possibility of a non-conforming implementation. We argue that validation of output models is still necessary: a faulty model transformation, however correctly implemented, will not improve quality. In addition, metamodel-based approaches (based, e.g., on OCL specification) allow for a higher expressivity of constraints to be validated than, e.g., approaches based on XML-Schema.

A straightforward solution consists in devising a translator from the language in which the transformation was certified ($^+$ CAL or xEMOF) into Java, and certifying this translator. Assuming that verification of a (manually or automatically derived) Java implementation is required, the OCL invariants contained in the metamodel can still be reused by translating them to JML [BCC⁺05] as discussed in [Ham04]. As with all source-code level verification approaches, a larger state space has to be explored, thus reaching practical limits more quickly than for the model-level counterpart as a result of the faithful representation

```

algorithm test {
  variables
    g ∈ VertexGraph; root ∈ g.node;
    alloc = [v ∈ g.node ↦ MemnodeForVertex[v]];
    pointsTo = [v ∈ g.node ↦ SetToSequence(Targets[v, g]);
    current = root; next = Null; prev = Null; i = 1; backref = Null; bQuit = FALSE;
  {while(¬bQuit){
    \ * go down the leftmost branch
    while((current ≠ Null) ∧ (alloc[current].markBit = FALSE)){
      alloc[current].markBit := TRUE; alloc[current].flag := 1;
      if(Len(pointsTo[current]) > 0){
        i := alloc[current].flag; next := pointsTo[current][i];
        pointsTo[current][i] := prev; prev := current;
        alloc[current].flag := alloc[current].flag + 1; current := next;
      };
    }; \ * end of while current
    \ * retreat, all objects pointed from current have been visited
    while(∧ (prev ≠ Null)
      ∧ (alloc[prev].flag = Len(pointsTo[prev]) + 1)){
      i := Len(pointsTo[prev]); next := pointsTo[prev][i];
      pointsTo[prev][i] := current; current := prev; prev := next;
    }; \ * end of while prev
    \ * retreated back to the starting point
    if(prev = Null){bQuit := TRUE;};
    if(¬bQuit){ \ * visit subgraph to the right of prev
      i := alloc[prev].flag - 1; backref := pointsTo[prev][i];
      pointsTo[prev][i] := current; next := pointsTo[prev][alloc[prev].flag];
      pointsTo[prev][alloc[prev].flag] := backref;
      current := next; alloc[prev].flag := alloc[prev].flag + 1;
    };
  };
  assert(Topology(pointsTo) = g.edge);
  assert(ReachableFrom[g, root] = {v ∈ g.node : alloc[v].markBit});
}}

```

Figure 3: Schorr-Waite expressed in $^+$ CAL. Preconditions concern valid input (not shown), post-conditions (see the assert statements) verify certain topological invariants.

of Java Virtual Machine abstractions. For example, expressions in Java may have side-effects while OCL expressions are guaranteed to be read-only. Adopting EMOF models as the only mechanism to define state allows us to consider only those state evolutions allowed by EMOF, reducing the state space to explore. This is in tune with the principle of reasoning at the highest-level of abstraction possible, because it's more efficient.

The application of TLA $^+$ has been investigated also in other software engineering contexts. In the field of enterprise software architectures, model-checking of web service protocols is reported in [JLLV04]. An Eclipse-based text editor to support editing $^+$ CAL and TLA $^+$ specifications is described in [Zam05].

Expressing behavior at the level of object-oriented models is also the aim of approaches under the Executable UML umbrella [RFW $^+$ 04]. Most products in this category target the embedded systems or telecommunications market and are heavily focused on state-charts. Minor semantic differences between representative statechart dialects are analyzed

in detail in [CD05], as part of the UML 2.0 Semantics Project.

Another formalization of OCL with tool support for verification is KeY [ABB⁺03] which targets different verification use cases from the ones addressed in this paper. Its execution language is JavaCard, with both JML and a dedicated Dynamic Logic as verification backends. KeY is used in tandem with a commercial UML tool.

Brucker et. al. [BW06a] describe in detail a tool to transform UML+OCL into a formalization processable by a theorem prover. The same team has also mechanized a Hoare-calculus for an idealized object-oriented programming language [BW06b]. In principle, both tools can evolve into an integrated proof environment for object-oriented programs. As of now, verification based on interactive theorem-provers is not fully automatic (that's where the *interactive* comes in), requiring assistance from the user who has to understand the underlying logic and the deduction rules. Once a language-processing algorithm has been formulated in an imperative language amenable for Hoare-analysis, those verification conditions that cannot be automatically derived by the tool have to be specified manually. Proof tactics are then to be applied (automatically or assisted by the user) to discharge the verification conditions and therefore the (Hoare) pre- and postconditions for the algorithm as a whole. There is as of now no similar Hoare-calculus with theorem prover support for ⁺CAL. ⁺CAL is translated to TLA⁺, The Temporal Logic of Actions [Lam02]. This logic is examined by Merz in [Mer03].

5 Conclusions and Further Work

Given the remarkable progress during the last decade in the areas of model-checking and in anchoring the semantics of metamodeling, there is no reason preventing combining their strengths to increase the reliability of model compilers. Our findings confirm that language metamodeling techniques contribute not only to the productivity of MDS D but also to its quality. Our prototype aims at enabling the interchange of standard metamodels and certified transformations within the software engineering community, reaping the benefits of network effects. Integrated model-driven toolchains for enterprise-scale projects involve metamodels for several languages, whose development costs would be prohibitive if done from scratch and in isolation by separate teams. We foresee the existence of peer-reviewed, public repositories of machine-checked metamodels and transformations in the near future.

After certifying properties about well-formedness, a follow-up problem is *compiler correctness* [BGL05], i.e. answering whether the meaning of a high level specification is preserved in its compiled formulation. For example, whether an OCL query retrieves the same resultset as its SQL'92 formulation. In the example, the formal problem consists in determining whether the membership conditions defined over their respective universes by the OCL and the SQL'92 queries are logically equivalent. Research in this field for model compilers targeting an enterprise-class software architecture is at its early stages.

Our work on certifying model transformations will be continued with the development of extensions to the Eclipse Modeling infrastructure (supported by a 2006 Eclipse Innovation Grant). This project involves implementing transformations so that generated

Java code exhibits (a) efficient detection of broken OCL invariants, (b) reactive rules, and (c) statechart-related design patterns. There is an interplay between these research and development activities. The listed capabilities were motivated by our work on declarative specifications of behavior, as necessary for example when defining consistency-enforcing algorithms for use in DSL editors. In turn, the experience gained with case studies on model-checking transformations is directly applicable to the language design of the modeling infrastructure extensions mentioned above.

References

- [ABB⁺03] Wolfgang Ahrendt, Thomas Baar, Bernhard Beckert, Richard Bubel, Martin Giese, Reiner Hähnle, Wolfram Menzeand Wojciech Mostowski, Andreas Roth, Steffen Schlager, and Peter H. Schmitt. The KeY Tool. Technical report, Department of Computing Science, Chalmers University and Göteborg University, Göteborg, Sweden, 2003.
- [Ake04] D. H. Akehurst. Validating BPEL Specifications using OCL. Technical Report 15-04, University of Kent at Canterbury, August 2004.
- [AKRS06] 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.
- [APM03] G. Antoniol, M. Di Penta, and E. Merlo. YAAB (Yet Another AST Browser): Using OCL to Navigate ASTs. In *IWPC '03: Proceedings of the 11th IEEE International Workshop on Program Comprehension*, page 13, Washington, DC, USA, 2003. IEEE Computer Society.
- [AS06] C. Amelunxen and A. Schürr. On OCL as part of the metamodeling framework MOFLON. In *6th OCL Workshop at the UML/ModelS Conference*, 2006.
- [BB06] Fabian Büttner and Hanna Bauerdick. Realizing UML Model Transformations with USE. In Dan Chiorean, Birgit Demuth, Martin Gogolla, and Jos Warmer, editors, *UML/ModelS Workshop on OCL (OCLApps'2006)*, pages 96–110. Technical University of Dresden, Technical Report TUD-FI06, 2006.
- [BCC⁺05] Lilian Burdy, Yoonsik Cheon, David Cok, Michael D. Ernst, Joe Kiniry, Gary T. Leavens, K. Rustan M. Leino, and Erik Poll. An overview of JML tools and applications. *Software Tools for Technology Transfer*, 7(3):212–232, June 2005.
- [BDW06] Achim D. Brucker, Jürgen Doser, and Burkhart Wolff. A Model Transformation Semantics and Analysis Methodology for SecureUML. Technical Report 524, ETH Zürich, 2006.
- [BGL05] Jan Olaf Blech, Sabine Glesner, and Johannes Leitner. Formal Verification of Java Code Generation from UML Models. Fujaba Days, september 2005.
- [BKS02] Bernhard Beckert, Uwe Keller, and Peter H. Schmitt. Translating the Object Constraint Language into First-order Predicate Logic. In *Proceedings, VERIFY, Workshop at Federated Logic Conferences (FLoC), Copenhagen, Denmark, 2002*. Available at i12www.ira.uka.de/key/doc/2002/BeckertKellerSchmitt02.ps.gz.

- [BSM⁺03] Frank Budinsky, David Steinberg, Ed Merks, Raymond Ellersick, and Timothy J. Grose. *Eclipse Modeling Framework*. Addison-Wesley Professional, Boston, MA, USA, 2003.
- [BW06a] Achim D. Brucker and Burkhard Wolff. The HOL-OCL Book. Technical Report 525, ETH Zürich, 2006.
- [BW06b] Achim D. Brucker and Burkhard Wolff. A Package for Extensible Object-Oriented Data Models with an Application to IMP++. In Abhik Roychoudhury and Zijiang Yang, editors, *International Workshop on Software Verification and Validation (SVV 2006)*, Computing Research Repository (CoRR). Seattle, USA, August 2006.
- [CD05] Michelle L. Crane and Jürgen Dingel. UML Vs. Classical Vs. Rhapsody Statecharts: Not All Models Are Created Equal. In Lionel C. Briand and Clay Williams, editors, *MoDELS*, volume 3713 of *Lecture Notes in Computer Science*, pages 97–112. Springer, 2005.
- [Cer05] Stefano Ceri. Process Modeling in Web Applications. In Robert Meersman, Zahir Tari, Mohand-Said Hacid, John Mylopoulos, Barbara Pernici, Özalp Babaoglu, Hans-Arno Jacobsen, Joseph P. Loyall, Michael Kifer, and Stefano Spaccapietra, editors, *OTM Conferences (1)*, volume 3760 of *Lecture Notes in Computer Science*, page 20. Springer, 2005.
- [EEHT05] Karsten Ehrig, Claudia Ermel, Stefan Hänsgen, and Gabriele Taentzer. Generation of visual editors as eclipse plug-ins. In *ASE '05: Proceedings of the 20th IEEE/ACM international Conference on Automated software engineering*, pages 134–143, New York, NY, USA, 2005. ACM Press.
- [Gar] M. Garcia. Accompanying materials to this paper
<http://www.sts.tu-harburg.de/~mi.garcia/pubs/2007/se2007>.
- [Gar06] M. Garcia. Formalizing the well-formedness rules of EJB3QL in UML + OCL. In *3rd Intl Workshop on Metamodels, Schemas, Grammars, and Ontologies for Reverse Engineering, at the UML/MoDELS Conference*, 2006.
- [Gol05] Bob Goldberg. The DASL Language: Programmer’s Guide and Reference Manual, TR-2005-128. Technical report, Sun Microsystems Research Labs, 2005.
- [GVJH98] Erich Gamma, John Vlissides, Ralph Johnson, and Richard Helm. *Design Patterns CD: Elements of Reusable Object-Oriented Software, (CD-ROM)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1998.
- [Ham04] Ali Hamie. Translating the Object Constraint Language into the Java Modelling Language. In *SAC '04: Proceedings of the 2004 ACM symposium on Applied computing*, pages 1531–1535, New York, NY, USA, 2004. ACM Press.
- [HM05] Thierry Hubert and Claude Marche. A case study of C source code verification: the Schorr-Waite algorithm. In *SEFM '05: Proceedings of the Third IEEE International Conference on Software Engineering and Formal Methods*, pages 190–199, Washington, DC, USA, 2005. IEEE Computer Society.
- [HZS05] Shan Shan Huang, David Zook, and Yannis Smaragdakis. Statically Safe Program Generation with SafeGen. In Robert Glück and Michael R. Lowry, editors, *GPCE*, volume 3676 of *Lecture Notes in Computer Science*, pages 309–326. Springer, 2005.
- [Jac06] Daniel Jackson. *Software Abstractions: Logic, Language, and Analysis*. The MIT Press, 2006.

- [JBK06] Frédéric Jouault, Jean Bézivin, and Ivan Kurtev. TCS: a DSL for the Specification of Textual Concrete Syntaxes in Model Engineering. In *GPCE'06: Proceedings of the Fifth International Conference on Generative programming and Component Engineering*, 2006. To appear.
- [JK06] Frédéric Jouault and Ivan Kurtev. On the Architectural Alignment of ATL and QVT. In *Proceedings of the 2006 ACM Symposium on Applied Computing (SAC 06)*, pages 1188–1195, Dijon, France, 2006. ACM Press.
- [JLLV04] James E. Johnson, David E. Langworthy, Leslie Lamport, and Friedrich H. Vogt. Formal Specification of a Web Services Protocol. *Electr. Notes Theor. Comput. Sci.*, 105:147–158, 2004.
- [Lam02] Leslie Lamport. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2002.
- [Lam06a] L. Lamport. Checking a Multithreaded Algorithm with +CAL. In *Proc. 20th International Symposium on Distributed Computing DISC 06*, 2006.
- [Lam06b] Leslie Lamport. The +CAL Algorithm Language. *Submitted for publication*, 2006. Available at <http://research.microsoft.com/users/lamport/pubs/pluscal.pdf>.
- [Mer03] Stephan Merz. On the Logic of TLA⁺. *Computers and Informatics*, 22:351–379, 2003.
- [MMS98] Bertrand Meyer, Christine Mingins, and Heinz Schmidt. Providing Trusted Components to the Industry. *Computer*, 31(5):104–105, 1998.
- [MN05] Farhad Mehta and Tobias Nipkow. Proving pointer programs in higher-order logic. *Inf. Comput.*, 199(1-2):200–227, 2005.
- [OMG05] Object Management Group. MOF QVT Final Adopted Specification, formal/05-11-01, Nov 2005.
- [OMG06] Object Management Group. Meta Object Facility (MOF) Core Specification, formal/06-01-01, Jan 2006.
- [RFW⁺04] Chris Raistrick, Paul Francis, John Wright, Colin Carter, and Ian Wilkie. *Model Driven Architecture with Executable UML*. Cambridge University Press, Cambridge, UK, 2004.
- [WK03] Jos Warmer and Anneke Kleppe. *The Object Constraint Language: Getting Your Models Ready for MDA*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2003.
- [WKC06] Junhua Wang, Soon-Kyeong Kim, and David Carrington. Verifying Metamodel Coverage of Model Transformations. In *ASWEC '06: Proceedings of the Australian Software Engineering Conference (ASWEC'06)*, pages 270–282, Washington, DC, USA, 2006. IEEE Computer Society.
- [Zam05] Boris Gruschko; Friedrich H. Vogt; Simon Zambrovski. The Use of TLA+ and Model Checking Tools in the Eclipse Environment. In *2nd International Workshop on Web Services and Formal Methods*, Versailles, France, 9 2005.