

An Algebraic Semantics for MOF

Artur Boronat¹ and José Meseguer²

¹ Department of Computer Science, University of Leicester
aboronat@le.ac.uk

² Department of Computer Science, University of Illinois at Urbana-Champaign
meseguer@uiuc.edu

Abstract. In model-driven development, software artifacts are represented as models in order to improve productivity, quality, and cost effectiveness. In this area, the Meta-Object Facility (MOF) standard plays a crucial role as a generic framework within which a wide range of modeling languages can be defined. The MOF standard aims at offering a good basis for model-driven development, providing some of the building concepts that are needed: what is a model, what is a metamodel, what is reflection in the MOF framework, and so on. However, most of these concepts are not yet fully formally defined in the current MOF standard. In this paper we define a reflective, algebraic, executable framework for precise metamodeling based on membership equational logic (MEL) that supports the MOF standard. Our formal framework provides a formal semantics of the following notions: *metamodel*, *model*, and *conformance* of a model to its metamodel. Furthermore, by using the Maude language, which directly supports MEL specifications, this formal semantics is *executable*. This executable semantics has been integrated within the Eclipse Modeling Framework as a plugin tool called MOMENT2. In this way, formal analyses, such as semantic consistency checks, model checking of invariants and LTL model checking, become available within Eclipse to provide formal support for model-driven development processes.

Key words: MOF, model-driven development, membership equational logic, metamodeling semantics, reflection, formal analysis.

1 Introduction

Model-driven development is a software engineering field in which software artifacts are represented as models in order to improve productivity, quality, and cost-effectiveness. Models provide a more abstract description of a software artifact than the final code of the application. The Meta-Object Facility (MOF) standard [45] offers a generic framework in which the abstract syntax of different modeling languages can be defined. This is done by specifying within MOF different metamodels for different modeling languages. Models in a modeling language are then conforming instances of their corresponding metamodel. The MOF standard aims at offering a good basis for model-driven development, providing some of the building concepts that are needed: what is a model, what is

a metamodel, what is reflection in the MOF framework, and so on. However, most of these concepts are not yet fully formally defined in the current MOF standard. This is, in part, due to the fact that metamodels can only be defined as data in the MOF framework.

In this paper, we define a reflective, algebraic, executable framework for precise metamodeling that supports the MOF standard. Our formal framework provides a formal semantics of the following notions: *metamodel*, *model*, and *conformance* of a model to its metamodel. We clearly distinguish the different roles that the notion of *metamodel* usually plays in the literature: as *data*, as *type*, and as *theory*. In addition, we introduce two new notions: (i) *metamodel realization*, referring to the mathematical representation of a metamodel; and (ii) *model type*, allowing models to be considered as first-class citizens. In particular, our executable algebraic semantics for MOF generates, in an automatic way, the algebraic semantics of any MOF metamodel by formalizing MOF metamodels as MEL theories. Furthermore, such MEL theories are executable by rewriting. This means that useful MEL deduction capabilities become available as decision procedures to check properties about the models of a metamodel so formalized; and that all this can be achieved by executing the MEL theory formalizing a metamodel in an algebraic language such as Maude [22]. In this way, the executable formal semantics of a metamodel can be used to automatically analyze the conformance of its model instances. Such model instances are characterized algebraically as terms modulo structural axioms of associativity, commutativity and identity, and have an equivalent topological characterization as graphs. This makes the formal semantics particularly useful, since models can be directly manipulated as graphs in their term-modulo-axioms formal representation. Furthermore, as we explain below, it makes possible a number of very useful model management and model analysis tasks.

Our framework provides not only an algebraic semantics, but also an executable environment called MOMENT2 [43], that is plugged into the Eclipse Modeling Framework (EMF) [25] and that constitutes the kernel of an algebraic model management framework supporting model transformations and formal analysis techniques. In this work, we illustrate the basic principles on which MOMENT2 is based by showing how the executable formalization of MOF metamodels as MEL theories provides automated, semantics-based support for very useful model-based tasks, such as the definition of domain specific languages, model transformations, model traceability and model management operators; and formal analysis techniques, such as reachability analysis and LTL model checking of model-based systems.

Philosophically, we can view this work as exploiting for the area of software modeling languages the kinds of benefits that the rewriting logic semantics project [42] has already demonstrated in the areas of programming languages (see [41,42,54] and references there), and of formal specification and verification languages (see, e.g., [23,55,19]). What is common to all these efforts is the use of a flexible logical framework (either rewriting logic [38], or, as done in this paper, its MEL equational sublogic [40]) to formalize the semantics of either a

programming language, a formal specification language, or, as done in this work, a wide range of software modeling languages. By exploiting the fact that the logical framework’s logic is executable and has a high-performance implementation, one obtains much more than just precise semantic definitions: one obtains useful (and quite efficient) semantics-based tools, such as programming language interpreters and model checkers, theorem provers of various kinds, and, in this paper, a model management tool like MOMENT2 that is metamodel-generic and can support a wide range of model management and model analysis tasks.

The paper is structured as follows: Section 2 briefly describes the underlying formal background; Section 3 identifies important concepts that are not formally defined in the MOF standard and are usually left unspecified in most of the MOF implementations; Section 4 presents the foundations of our algebraic framework, indicating how the algebraic semantics of MOF metamodels is defined; Section 5 illustrates the application of our framework in different model-driven development scenarios; Section 6 discusses related work; and Section 7 summarizes the main contributions of this work and discusses future work.

2 Membership Equational Logic and Maude

The logical framework in which we give an algebraic semantics to MOF is membership equational logic (MEL) [40]. A key feature of MEL is its strong support for types (called sorts), subtypes, and partiality; and the closely-related feature that types can be very expressive. In MEL, membership in a type is not just a syntactic matter, but may depend on the satisfaction of *semantic conditions*. This MEL feature is exploited very heavily in our semantics of MOF, where the set of models that conform to a given metamodel is precisely characterized by a type satisfying suitable semantic conditions. Another key feature of MEL is that it has *initial algebras* [40]; therefore, our MOF semantics is an initial algebra semantics. Furthermore, under very reasonable assumptions MEL theories are *executable* by rewriting. Using Maude’s high-performance implementation of MEL, this means that theories satisfying such executability conditions become declarative programs, and that various MEL deduction tasks, including checking equality and checking membership in a type, become decision procedures.

2.1 The Syntax and Semantics of MEL

A MEL *signature* is a triple (K, Σ, S) (just Σ in the following), with K a set of *kinds*, $\Sigma = \{\Sigma_{w,k}\}_{(w,k) \in K^* \times K}$ a many-kinded signature of function symbols, and $S = \{S_k\}_{k \in K}$ a K -kinded family of disjoint sets of sorts. The kind of a sort s is denoted by $[s]$. Intuitively, membership in a sort is exactly *definedness*; whereas membership in just a kind without membership in any of the sorts of that kind means *undefinedness* or *error*. For example, let *Numeric* be a kind having sorts *Nat*, *Int*, and *Rat*. Then, in a suitable MEL specification of the number hierarchy we may have expressions like $4 - 7$ of sort *Int*, and $2/7$ of sort

Rat. But the expression $2/0$ has kind *Numeric* but has no sort and is therefore interpreted as an undefined or error expression.

A MEL Σ -algebra A contains a set A_k for each kind $k \in K$, a function $A_f : A_{k_1} \times \cdots \times A_{k_n} \rightarrow A_k$ for each operator $f \in \Sigma_{k_1 \dots k_n, k}$ and a subset $A_s \subseteq A_k$ for each sort $s \in S_k$, with the meaning that the elements in sorts are well-defined, while elements without a sort are *errors*. $T_{\Sigma, k}$ and $T_{\Sigma}(X)_k$ denote, respectively, the set of ground Σ -terms with kind k and of Σ -terms with kind k over variables in X , where $X = \{x_1 : k_1, \dots, x_n : k_n\}$ is a set of kinded variables.

Given a MEL signature Σ , *atomic formulae* have either the form $t = t'$ (Σ -equation) or $t : s$ (Σ -membership) with $t, t' \in T_{\Sigma}(X)_k$ and $s \in S_k$; and Σ -sentences are conditional formulae of the form $(\forall X) \varphi$ *if* $\bigwedge_i p_i = q_i \wedge \bigwedge_j w_j : s_j$, where φ is either a Σ -equation or a Σ -membership, and all the variables in φ , p_i , q_i , and w_j are in X . The novel feature with respect to other equational logics is the support for conditional memberships of the form $(\forall X) t : s$ *if* $\bigwedge_i p_i = q_i \wedge \bigwedge_j w_j : s_j$. This makes membership in a sort not just a syntactic matter, but a semantic one, since the semantic conditions $\bigwedge_i p_i = q_i \wedge \bigwedge_j w_j : s_j$ must be satisfied.

A MEL theory is a pair (Σ, E) with Σ a MEL signature and E a set of Σ -sentences. The paper [40] gives a detailed presentation of (Σ, E) -algebras, sound and complete deduction rules, and initial and free algebras. In particular, given a MEL theory (Σ, E) , its initial algebra is denoted $T_{(\Sigma/E)}$; its elements are E -equivalence classes of ground terms in T_{Σ} .

Order-sorted notation $s_1 < s_2$ can be used to abbreviate the conditional membership $(\forall x : k) x : s_2$ *if* $x : s_1$. Similarly, an operator declaration $f : s_1 \times \cdots \times s_n \rightarrow s$ corresponds to declaring f at the kind level and giving the membership axiom $(\forall x_1 : k_1, \dots, x_n : k_n) f(x_1, \dots, x_n) : s$ *if* $\bigwedge_{1 \leq i \leq n} x_i : s_i$. We write $(\forall x_1 : s_1, \dots, x_n : s_n) t = t'$ in place of $(\forall x_1 : k_1, \dots, x_n : k_n) t = t'$ *if* $\bigwedge_{1 \leq i \leq n} x_i : s_i$.

We can use order-sorted notation as syntactic sugar to present a MEL theory (Σ, E) in a more readable form as a tuple $(S, <, \Sigma, E_0 \cup A)$ where: (i) S is the set of sorts; (ii) $<$ is the subsort inclusions, so that there is an implicit kind associated to each connected component in the poset of sorts $(S, <)$; (iii) Σ is given as an order-sorted signature with possibly overloaded operator declarations $f : s_1 \times \cdots \times s_n \rightarrow s$ as described above; and (iv) the set E of (possibly conditional) equations and memberships is quantified with variables having specific sorts (instead than with variables having specific kinds) in the sugared fashion described above; furthermore, E is decomposed as a disjoint union $E = E_0 \cup A$, where A is a collection of “structural” axioms such as associativity, commutativity, and identity. Any theory $(S, <, \Sigma, E_0 \cup A)$ can then be desugared into a standard MEL theory (Σ, E) in the way explained above.

2.2 Executable MEL Theories and Maude

The point of the decomposition $E = E_0 \cup A$ is that, under appropriate executability requirements explained in [18,22], such as confluence, termination, and sort-decreasingness modulo A , a MEL theory $(S, <, \Sigma, E_0 \cup A)$ becomes *executable*

by rewriting with the equations and memberships E_0 modulo the structural axioms A . Furthermore, the initial algebra $T_{(\Sigma/E)}$ then becomes isomorphic to the *canonical term algebra* $Can_{\Sigma/E_0,A}$ whose elements are A -equivalence classes of ground Σ -terms that cannot be further simplified by the equations and memberships in E_0 .

An important consequence of $(S, <, \Sigma, E_0 \cup A)$ satisfying the confluence, termination and sort-decreasingness requirements is that both term equality and term membership become *decidable* by rewriting [18]. Furthermore, MEL theories satisfying these executability properties can be executed in high-performance languages such as Maude [22] that support execution of MEL theories modulo any combination of associativity and/or commutativity and/or identity axioms.

The syntax of Maude follows very closely the mathematical syntax for MEL described above, including support for order-sorted notation. In the rest of the paper we will at times present fragments of MEL specifications in Maude syntax. Such syntax is essentially self-explanatory, but we give here a few explanations to help the reader. First of all, basic notions such as sorts, subsorts, operators, equations, and memberships are declared with respective keywords `sort`, `subsort`, `op`, `eq` (or `ceq` for conditional equations), and `mb` (or `cmb` for conditional memberships). Second, associativity and/or commutativity and/or identity axioms are not declared explicitly as equations, but are instead declared as attributes of the binary operator enjoying them in the corresponding `op` declaration with the `assoc`, `comm`, and `id`: keywords. Third, the syntax of an operator need not be prefix syntax: it can also be *user-definable* “mixfix” syntax, where the argument places are indicated by underbars. For example, a user can define an if-then-else operator with syntax `if_then_else.fi`.

3 MOF and its Semantic Issues

In this section, we give an informal description of the MOF standard by describing the MOF architecture and the main concepts in the MOF metamodel. We also discuss several semantic issues that are then addressed by the formal algebraic semantics presented in subsequent sections.

3.1 The MOF Modeling Framework

MOF is a semiformal approach to define modeling languages. It provides a four-level hierarchy, with levels M0, M1, M2 and M3. Each entity M at level i represents a *model*³. M is constituted by a collection of *typed* elements e , which can be simple data values or objects. Types T for the elements e that constitute a model M at level i are defined as collections of elements in a *metamodel* \mathcal{M} at level $i + 1$. There are two different kinds of types T : data types and object types. We indicate that an element e is *typed* with a type T as $e : T$. This relation is

³ In the MOF framework, the concept of a *model* M is conceptually specialized depending on the specific metalevel, in which a model is located: *model* at level M1, *metamodel* at level M2 and *meta-metamodel* at level M3; as shown below.

specialized for simple data values as the *isValueOf* relation, and for objects as the *isInstanceOf* relation. In addition, a model M is said to *conform* to a metamodel \mathcal{M} [8], when the elements that constitute the model are properly typed. We call this relation *conformance relation* between a model M and a metamodel \mathcal{M} and we denote it by $M : \mathcal{M}$.

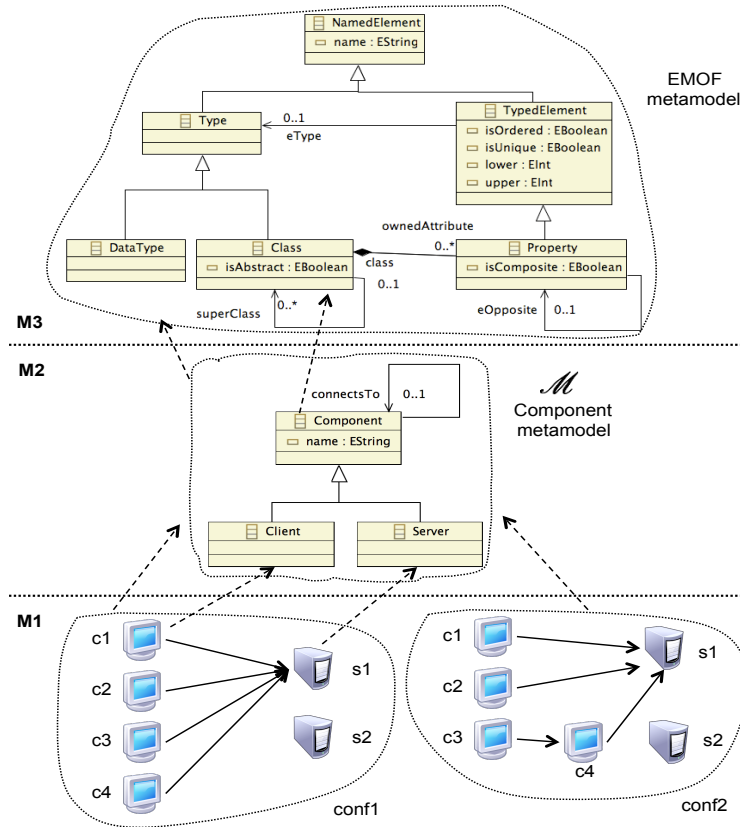


Fig. 1. The MOF framework.

Fig. 1 illustrates models at each level M1-M3 of the MOF framework. Each model is encircled by a boundary and tagged with a name. For example, `conf1` at level 1, which is a model corresponding to a configuration of components. The *isInstanceOf* relation between objects of a model M and the metarepresentation of object types in a metamodel \mathcal{M} , and the *conformance* relation between a model M and the corresponding metamodel \mathcal{M} are depicted with dashed arrows. We consider levels M1–M3 out of the MOF hierarchy in this work, which are:

M1 level. This level contains metarepresentations of *models*. A model is a collection of data elements that describe the elements of some physical, abstract

or hypothetical reality by using a well-defined language. In addition, a model is suitable for computer-based interpretation, so that development tasks can be automated. For example, the model `conf1` in Fig. 1 defines a configuration of client/server components describing how they are interconnected. Each icon corresponds to a component indicating whether it is a client or a server. We use this graphical representation to enhance the visualization of a model. In addition, current technology permits using both textual and graphical concrete syntax for MOF models enhancing the definition of domain-specific languages for end-users.

M2 level. This level contains metarepresentations of *metamodels*. A metamodel is a model specifying the abstract syntax of a modeling language. As an example, we take a metamodel for defining configurations of components as shown in Fig. 1 in UML notation. The types of a configuration are called *Component*, *Server* and *Client*. The *Component* object type is *abstract* so that it cannot be instantiated. This object type is specialized into a *Client* and a *Server* object types. The model `conf1` at level 1 conforms to the component metamodel \mathcal{M} at level 2.

M3 level. An entity at level 3 is the metarepresentation of a *meta-metamodel*. A meta-metamodel specifies a *modeling framework*, which could also be called a *modeling space*. In MOF, there is only one such meta-metamodel, called the MOF meta-metamodel (usually also called the MOF metamodel). Within the MOF standard, we focus on the *Essential MOF* (EMOF) specification that describes the meta-metamodel, which appears simplified in Fig. 1. In the MOF modeling framework one can define many different metamodels. Such metamodels, when represented as data, must conform to the MOF meta-metamodel. In particular, the component metamodel conforms to the MOF meta-metamodel. Other metamodels that are likewise specified in MOF are the UML2 metamodel to define UML models, the OWL metamodel to define ontologies [35], the AADL metamodel to specify real-time distributed computer systems [51], and so on. The fact that all these metamodels are specified within the single MOF framework greatly facilitates systematic model/metamodel interchange and integration.

3.2 Discussion and Open Problems

At present, important MOF concepts such as those of metamodel, model and conformance relation do not have an explicit, *syntactically characterizable* status in their data versions. For example, we can syntactically characterize the correctness of the data elements of a metamodel \mathcal{M} , but there is no explicit type that permits defining \mathcal{M} as a well-characterized value, which we call a *model type*. In addition, in the MOF standard and in current MOF-like modeling environments, such as Eclipse Modeling Framework or MS DSL tools, a metamodel \mathcal{M} does not have a precise *mathematical* status. Instead, at best, a metamodel \mathcal{M} is realized as a program in a conventional language as, for example, the Java code that is generated for a metamodel \mathcal{M} in EMF. This informal implementation corresponds to what we call a *metamodel realization*. In these

modeling environments, the *conformance relation* between a model definition M and its corresponding metamodel \mathcal{M} is checked by means of indirect techniques based on XML document validation or on tool-specific implementations in OO programming languages. Therefore, metamodels \mathcal{M} and models M cannot be explicitly characterized as first-class entities in their data versions, and the semantics of the *conformance relation* remains formally unspecified. This is due to the lack of a suitable reflective formal framework in which software artifacts, and not just their metarepresentations, can acquire a formal semantics.

In this work, we formalize the notions of: (i) model type, (ii) metamodel realization, and (iii) conformance relation, by means of a reflective semantics that associates a mathematical metamodel realization to each metamodel \mathcal{M} in MOF.

4 An Algebraic Semantics for MOF

The practical usefulness of a formal semantics for a language is that it provides a rigorous standard that can be used to judge the correctness of an implementation. For example, if a programming language lacks a formal semantics, compiler writers may interpret the informal semantics of the language in different ways, resulting in possibly inconsistent and diverging implementations. For MOF, given its genericity, the need for a formal semantics that can serve as a rigorous standard for any implementation is even more pressing, since many different modeling languages rely on the correctness of the MOF infrastructure. In this section, we propose an algebraic, mathematical semantics for MOF in membership equational logic (MEL).

4.1 A High-Level View of the MOF Algebraic Semantics

A metamodel \mathcal{M} describes a metamodel realization that contains a model type. What this metamodel describes is, of course, a *set* of models. We call this the *extensional semantics* of \mathcal{M} , and denote this semantics by $\llbracket \mathcal{M} \rrbracket$. Recall that we use the notation $M : \mathcal{M}$ for the conformance relation. Using this notation, the extensional semantics can be informally defined as follows:

$$\llbracket \mathcal{M} \rrbracket = \{M \mid M : \mathcal{M}\}.$$

We make the informal MOF semantics just described mathematically precise in terms of the *initial algebra semantics* of MEL. As already mentioned in Section 2, a MEL specification (Σ, E) has an associated initial algebra $T_{(\Sigma, E)}$. We call $T_{(\Sigma, E)}$ the *initial algebra semantics* of (Σ, E) , and write

$$\llbracket (\Sigma, E) \rrbracket_{IAS} = T_{(\Sigma, E)}.$$

Let $\llbracket \text{MOF} \rrbracket$ denote the set of all MOF metamodels \mathcal{M} , and let *SpecMEL* denote the set of all MEL specifications. The reason why we define $\llbracket \text{MOF} \rrbracket$ as a set of metamodels \mathcal{M} , instead than as a set of model types $\llbracket \mathcal{M} \rrbracket$ is because,

as already mentioned, the mathematical status of \mathcal{M} is, as yet, undefined, and is precisely one of the questions to be settled by a mathematical semantics. Instead, well-formed metamodels \mathcal{M} are data structures that can be syntactically characterized in a fully formal way. Therefore, the set $\llbracket \text{MOF} \rrbracket$, thus understood, is a well-defined mathematical entity. Our algebraic semantics is then defined as a function

$$\mathbb{A} : \llbracket \text{MOF} \rrbracket \longrightarrow \text{SpecMEL}$$

that associates to each MOF metamodel \mathcal{M} a corresponding MEL specification $\mathbb{A}(\mathcal{M})$. Our informal semantics $\llbracket \mathcal{M} \rrbracket$ is now made mathematically precise. Recall that any MEL signature Σ has an associated set S of sorts. Therefore, in the initial algebra $T_{(\Sigma, E)}$ each sort $s \in S$ has an associated set of elements $T_{(\Sigma, E), s}$. The key point is that in any MEL specification of the form $\mathbb{A}(\mathcal{M})$, there is always a sort called *Model*, whose data elements in the initial algebra are precisely the data representations of those models that conform to \mathcal{M} . That is, *Model* is the syntactical representation of the *model type* $\llbracket \mathcal{M} \rrbracket$ associated to a metamodel \mathcal{M} . Therefore, we can give a precise mathematical semantics to our informal MOF extensional semantics by the defining equation

$$\llbracket \mathcal{M} \rrbracket = T_{\mathbb{A}(\mathcal{M}), \text{Model}}.$$

Note that our algebraic semantics gives a precise mathematical meaning to the entities lacking such a precise meaning in the informal semantics, namely, the notions of: (i) model type $\llbracket \mathcal{M} \rrbracket$, (ii) metamodel realization $\mathbb{A}(\mathcal{M})$, and (iii) conformance relation $M : \mathcal{M}$. Specifically, we associate to a metamodel \mathcal{M} a precise mathematical object, namely, the MEL theory $\mathbb{A}(\mathcal{M})$, constituting its *metamodel realization*. The *structural conformance* relation between a model and its metamodel is then defined mathematically by the equivalence

$$M : \mathcal{M} \iff M \in T_{\mathbb{A}(\mathcal{M}), \text{Model}}.$$

4.2 Reflective Bootstrapping of the Algebraic Semantics

The algebraic semantics that we propose exploits the reflective features of both MOF and MEL. This allows a modular, stepwise approach in the definition of the semantic function \mathbb{A} . This has many advantages, both theoretically and in the practical realization of \mathbb{A} in the MOMENT2 tool.

The key observation about reflection in MOF is that the MOF meta-meta-model at level 3 is also a meta-model at level 2, which can be treated just as any other metamodel. In particular this means that

$$\text{MOF} \in \llbracket \text{MOF} \rrbracket.$$

It also means that our algebraic semantics function $\mathbb{A} : \llbracket \text{MOF} \rrbracket \longrightarrow \text{SpecMEL}$ applies in particular to MOF, that is, that there is a MEL theory $\mathbb{A}(\text{MOF})$.

The MEL theory $\mathbb{A}(\text{MOF})$ is enormously useful to bootstrap our algebraic semantics because of the following remarkable property. Since for any metamodel \mathcal{M} we have the identity $\llbracket \mathcal{M} \rrbracket = T_{\mathbb{A}(\mathcal{M}), \text{Model}}$, in particular for the metamamodel MOF we have the identity

$$\llbracket \text{MOF} \rrbracket = T_{\mathbb{A}(\text{MOF}), \text{Model}}.$$

This suggests the following bootstrapping strategy to define the semantic function $\mathbb{A} : \llbracket \text{MOF} \rrbracket \rightarrow \text{SpecMEL}$.

1. First define the MEL theory $\mathbb{A}(\text{MOF})$. This automatically gives us the domain $\llbracket \text{MOF} \rrbracket$ of our desired semantic function $\mathbb{A} : \llbracket \text{MOF} \rrbracket \rightarrow \text{SpecMEL}$ as the algebraic data type $\llbracket \text{MOF} \rrbracket = T_{\mathbb{A}(\text{MOF}), \text{Model}}$.
2. Once the domain $\llbracket \text{MOF} \rrbracket$ of \mathbb{A} is thus algebraically represented, proceed to give a recursive definition of the semantic function \mathbb{A} for any $\mathcal{M} \in \llbracket \text{MOF} \rrbracket$.

Step (1) is described in Section 4.3, and Step (2) in Section 4.4.

There is, however, a third very important step, namely, to also exploit reflection in the target logical framework MEL. Reflection in MEL means that there is a *universal* MEL theory $\mathcal{U} \in \text{SpecMEL}$ that can simulate the deduction of all finitary MEL theories, including its own deduction [24].

In particular, the universal theory \mathcal{U} has a sort *Module* that meta-represents all finitary MEL theories, including \mathcal{U} itself. This means that there is an algebraic data type $\overline{\text{SpecMEL}}$ whose elements are meta-representations of MEL theories, so that given any finitary MEL theory (Σ, E) , we have $(\Sigma, E) \in \text{SpecMEL}$ iff $(\overline{\Sigma}, \overline{E}) \in \overline{\text{SpecMEL}}$, where $(\overline{\Sigma}, \overline{E})$ is the meta-representation of (Σ, E) . Specifically, the algebraic data type $\overline{\text{SpecMEL}}$ has the following initial algebra semantics definition:

$$\overline{\text{SpecMEL}} = T_{\mathcal{U}, \text{Module}}.$$

All this means that we can take a third step in our bootstrapping process, namely, to realize our semantic function \mathbb{A} as an *equationally defined function*

$$\overline{\mathbb{A}} : \llbracket \text{MOF} \rrbracket \rightarrow \overline{\text{SpecMEL}}$$

mapping each metamodel \mathcal{M} to the metarepresentation $\overline{\mathbb{A}(\mathcal{M})}$ of the MEL theory $\mathbb{A}(\mathcal{M})$.

At the theoretical level, this third bootstrapping step means that the entire algebraic semantics of MOF can be defined *within* the semantic framework of MEL. At the practical level it also means that, because of the efficient support in Maude for key functionality of the universal MEL theory \mathcal{U} by means of Maude's **META-LEVEL** module, the function $\overline{\mathbb{A}}$ can be efficiently implemented in Maude. In fact, this is exactly how the core functionality of the MOMENT2 tool has been developed, as explained in Section 4.4.

In Sections 4.3 and 4.4 we give a high level summary of the definition of the \mathbb{A} semantic function, as achieved by the three reflective bootstrapping steps described above. A complete definition of \mathbb{A} is available in [11]⁴. In Section 4.5 we explain how \mathbb{A} is used in the EMF by means of MOMENT2.

⁴ In the dissertation, the \mathbb{A} function is denoted by $\text{reflect}_{\text{MOF}}$.

4.3 The MEL Theory $\mathbb{A}(\text{MOF})$

As already mentioned, we denote the metamodel that constitutes the meta-metamodel of the MOF framework by MOF. MOF is itself a MOF metamodel, since $\text{MOF} : \text{MOF}$. In this section we take the first bootstrapping step described in Section 4.2, namely, to define the MEL theory $\mathbb{A}(\text{MOF})$. That is, we first define \mathbb{A} for a *single* metamodel, namely MOF. The $\mathbb{A}(\text{MOF})$ theory defines the $[[\text{MOF}]]$ type as the set of metamodels \mathcal{M} , which can be viewed as both graphs and terms. This theory has been manually defined as a first step in the bootstrapping process needed to define the \mathbb{A} function in general.

The $\mathbb{A}(\text{MOF})$ theory provides the algebraic representation for both object types and model types for defining metamodels \mathcal{M} . In this theory, object types are used to describe a metamodel $\mathcal{M} : \text{MOF}$ as a collection of objects. Objects are defined by using the following sorts: `Oid` for object identifiers; `Cid` for class names; and `PropertySet` for multisets of comma-separated pairs of the form `(property : value)`, which represent property values. Objects in a metamodel \mathcal{M} are then syntactically characterized by means of an operator

$$\text{op } \langle _ : _ | _ \rangle : \text{Oid Cid PropertySet} \rightarrow \text{Object} .$$

These sorts, subsorts and operators are defined, in Maude notation, as follows:

```
sorts Oid Cid Property PropertySet Object .
subsort Property < PropertySet .
op noneProperty : -> PropertySet .
op _',_ : PropertySet PropertySet -> PropertySet
[assoc comm id: noneProperty] .
```

In the $\mathbb{A}(\text{MOF})$ theory, a metamodel \mathcal{M} that conforms to the meta-metamodel MOF, that is, such that $\mathcal{M} : \text{MOF}$, can be represented as a collection of objects by means of a term of sort `Model`. A term of sort `Model` is defined by means of the following constructors, in Maude notation:

```
op none : -> ObjCol .
op _> : ObjCol ObjCol -> ObjCol [assoc comm id: none] .
op <<_>> : ObjCol -> Model .
```

where the `Object` sort is a subsort of the `ObjCol` sort. That is, we first form a collection of objects of sort `ObjCol` using the associative and commutative multiset union operator $_>$ ⁵ and then we *wrap* the set of objects by using the `<<_>>` constructor to get the desired term of sort `Model`.

Each of the modeling primitives that constitute the MOF metamodel is specified in the $\mathbb{A}(\text{MOF})$ theory by means of sorts, subsorts and operators. As an example, we focus on the *Class* and *Property* object types of the MOF metamodel.

⁵ This binary operator symbol has empty syntax (juxtaposition).

Class *object type*. Object types are the central concept of MOF to model entities of the problem domain in metamodels. An object type is defined in a metamodel \mathcal{M} as a *Class* instance and a set of *Property* instances. The object type *Class* contains meta-properties like *name*, which indicates the name of the object type, *ownedAttribute*, which indicates the properties that belong to the *Class* instance, and *superClass*, which indicates that the object type is defined as a specialization of the object types that are referred to by means of this property. The *Class* object type is specified as a sort `Class`, such that `Class < Cid`, and a constant `Class : -> Class`. Each of the class properties is defined as a constructor for the sort `Property`. For example, to define the *name* property, we have the constructor `name' :_ : String -> Property`; to define the *isAbstract* property, we have the constructor `isAbstract' :_ : Bool -> Property`; to define the *ownedAttribute* property, we have the constructor `ownedAttribute' :_ : OrderedSet{Oid} -> Property`; and to define the *superClass* property, we have the constructor `superClass' :_ : Set{Oid} -> Property`.

In the component metamodel, the *Class* instance that defines the object type *Component* in the metamodel \mathcal{M} is defined as the term

```
< 'class0 : Class | name : "Component", isAbstract : true,
  ownedAttribute : .. >
```

where `'class0` is an object identifier.

Property *object type*. A model can be viewed as a graph where the collection of nodes is constituted by the collection of attributed objects of the model and the edges are defined by means of directed links between objects, defined as object-typed property values. Indeed, a model is an enriched graph where edges can also be defined as structural containments so that hierarchies of nested objects can also be represented.

A *Property* instance in a metamodel \mathcal{M} enables the definition of an *attribute* in an object or an *association end* between objects in a model definition $M : \mathcal{M}$, one level down in the MOF framework. A *Property* instance defines the type of the property, where the type is represented as a *DataType* instance (basic data values and enumerations) or as a *Class* instance (object types). Properties that are typed with datatypes are *attributes*, and properties that are typed with object types are *association ends* or *references*. An association end defines a unidirectional association between two classes. Bidirectional associations are defined by means of two *opposite* association ends. Composition associations can be defined by indicating that the association end that points to the composite class has its meta-property *isComposite* set to *true*. Other meta-properties, such as *lower*, *upper*, *isOrdered* and *isUnique*, constitute the multiplicity metadata of a specific property.

The constructors that permit defining objects of the *Property* object type are defined, in Maude notation, as follows⁶:

⁶ We represent the name of the *Property* class as `Prop` to avoid name collisions with the aforementioned sort *Property*.

```

sort Prop . subsort Prop < TypedElement .
op Prop : -> Prop .
op lower' :_ : Int -> Property .
op upper' :_ : Int -> Property .
op isOrdered' :_ : Bool -> Property .
op isUnique' :_ : Bool -> Property .
op opposite' :_ : Oid -> Property .
op opposite : -> Property .
    
```

The *Property* instance that defines the metaproperty *name* of the object type *Component* in the metamodel \mathcal{M} of the example is represented by the term

```
< 'prop : Prop | name : "name", type : 'String, class : 'class0 >.
```

where *'String* is the identifier of the object that represents the built-in String type. We have taken into account the modeling primitives that constitute the Essential MOF metamodel, including simple data types and enumeration types. A detailed specification is provided in [11]. The metamodel \mathcal{M} of the example is defined as a term of sort *Model* in the $\mathbb{A}(\text{MOF})$ theory in Fig. 2.

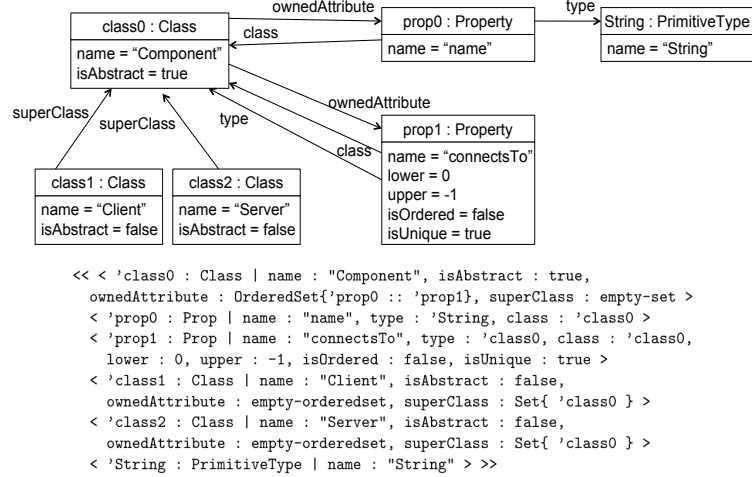


Fig. 2. Metamodel \mathcal{M} viewed as a graph and as a term.

4.4 Reflective Algebraic Semantics of MOF Metamodels

Once the $\mathbb{A}(\text{MOF})$ theory is defined, we focus on the *Model* sort in this theory, whose carrier in the initial algebra defines the $\llbracket \text{MOF} \rrbracket$ model type, i.e., the model type whose elements are metamodels:

$$\llbracket \text{MOF} \rrbracket = T_{\mathbb{A}(\text{MOF}), \text{Model}}.$$

Note that, since $\llbracket \text{MOF} \rrbracket$ is the set of all metamodels \mathcal{M} in MOF, this means that

$$\text{MOF} \in \llbracket \text{MOF} \rrbracket.$$

We then define the value of the function \mathbb{A} on *any* metamodel \mathcal{M} , such that $\mathcal{M} \in \llbracket \text{MOF} \rrbracket$, as its corresponding MEL theory $\mathbb{A}(\mathcal{M})$. This corresponds to the second bootstrapping step described in Section 4.2. Given a metamodel \mathcal{M} , the $\mathbb{A}(\mathcal{M})$ theory defines the $\llbracket \mathcal{M} \rrbracket$ semantics as the set of models M that are constituted by a collection of typed objects, which can be viewed as a graph where objects correspond to nodes, object attributes to node attributes and object-typed properties to edges. In addition, nodes can be nested by means of object-typed properties that are defined as *isComposite*. The metamodel \mathcal{M} of a model M corresponds to the type graph of the graph M , where object types constitute node types enabling node inheritance by means of class inheritance relationships.

In the $\mathbb{A}(\mathcal{M})$ theory, the algebraic notion of *object type* is generically given by means of the sort *Object*. Terms of sort *Object* are defined by means of the constructor

$$\langle _ : _ | _ \rangle : \text{Oid Cid PropertySet} \rightarrow \text{Object}$$

which is analogous to the constructor for objects that has been presented in the $\mathbb{A}(\text{MOF})$ theory.

Models $M : \mathcal{M}$ are given as collections of objects, which are instances of a specific object type OT. Object types are defined in a metamodel $\mathcal{M} : \text{MOF}$, as a collection of objects $\text{OT} \subseteq \mathcal{M}$. Defining the algebraic semantics of an object type involves the definition of the object identifiers and the properties that may be involved in the definition of a specific object in a model $M : \mathcal{M}$. Object type specialization relationships must be also taken into account. Therefore, we need to define the carrier of the sorts *Oid*, *Cid* and *PropertySet* for a specific object type definition.

Consider, for example, the metamodel \mathcal{M} in Fig. 1. In subsequent paragraphs, we use this example to obtain the theory that defines the *Component*, *Server* and *Client* object types.

Object Type Names. In the $\mathbb{A}(\mathcal{M})$ theory, each *Class* instance in \mathcal{M} is defined as a new sort and a constant, both of them with the name of the class. *Abstract classes* are defined as those that cannot be instantiated. The name of an abstract class C is *not* specified with a constant $C : C$, so that objects in a model M , such that $M : \mathcal{M}$, cannot have C as their type.

In the example of the component metamodel, the \mathbb{A} function generates a sort for the abstract object type *Component* in Maude notation as follows:

```
sort Component .
```

The concrete object types *Client* and *Server* are formalized as sorts and constants as follows:

```

sorts Client Server .
op Client : -> Client .
op Server : -> Server .

```

where the specialization of the *Component* class into the *Client* and *Server* classes is formalized as subsort relationships between the sorts of the corresponding classes:

```

subsorts Client Server < Component .

```

Object Type Properties. An object type is defined with a collection of *Property* instances describing its meta-properties. A *Property* instance is associated with a specific type t in the metamodel \mathcal{M} , which is defined as an object $t : \text{Type}$. Depending on the type t of a property, we can distinguish two kinds of properties:

- *Value-typed Properties or Attributes.* Properties of this kind are typed with *Data Type* instances, which can represent either a simple data type or an enumeration type. Value-typed properties define the attributes of nodes in a graph.
- *Object-typed Properties or Association Ends.* Properties of this kind are typed with object types, enabling the definition of unidirectional labelled edges in a graph.

The *type* meta-property together with the multiplicity metadata define a set of specific constraints on the acceptable values for the property type. These constraints are taken into account in the algebraic type that is assigned to the property by means of OCL collection types, as indicated in Table 1. In this table, the `Set{Oid}` sort corresponds to sets of identifiers that may be empty by means of the constant `empty-set`, while the `NeSet{Oid}` sort corresponds to sets of identifiers that cannot be empty, i.e., it excludes the constant `empty-set`. In the example, the *Component* object type is specified in the theory $\mathbb{A}(\mathcal{M})$, in Maude notation, as follows:

```

op name' : _ : String -> Property .
op connectsTo' : _ : Oid -> Property .
op connectsTo : -> Property .

```

where the constant `connectsTo` permits defining *unset connectsTo* properties in a model M since it is defined with lower bound 0 in the metamodel \mathcal{M} .

Object Type Specialization Relation. A *specialization* is a taxonomic relationship between two object types. This relationship specializes a general object type into a more specific one. In the example, we algebraically define the specialization relationship between the object types *Client* and *Component* by means of the subsort `Client < Component`. The supersorts of the resulting subsort hierarchy are defined as subsorts of the `Cid` sort, for object type name sorts and object identifier sorts, respectively.

Type	Lower Bound	Upper Bound	isOrdered	isUnique
\emptyset, Oid	0	1	-	-
$\text{Set}\{\text{Oid}\}$	0	*	false	true
$\text{OrderedSetOid}\}$	0	*	true	true
$\text{Bag}\{\text{Oid}\}$	0	*	false	false
$\text{Sequence}\{\text{Oid}\}$	0	*	true	false
Oid	1	1	-	-
$\text{NeSet}\{\text{Oid}\}$	1	*	false	true
$\text{NeOrderedSetOid}\}$	1	*	true	true
$\text{NeBag}\{\text{Oid}\}$	1	*	false	false
$\text{NeSequence}\{\text{Oid}\}$	1	*	true	false

Table 1. OCL types for encoding multiplicity constraints in properties.

Algebraic Semantics of Object Types. The algebraic semantics of an object type is then given by the set of all the objects that can be defined either as instances of the object type, i.e., a class, or as instances of any of its subtypes. The algebraic semantics of an object type definition OT is defined as follows:

$$\llbracket OT \rrbracket = \{o \mid o \in T_{\mathbb{A}(\mathcal{M}), \text{Object}} \wedge \text{class}(o) \in T_{\mathbb{A}(\mathcal{M}), \text{ClassSort}(OT)}\}.$$

where class is an operator that obtains the class name constant of a given object, and ClassSort is an operator that obtains the sort that corresponds to the object type in $\mathbb{A}(\mathcal{M})$. The isInstanceOf relation $o : OT$ is defined as follows:

$$o : OT \Leftrightarrow o \in \llbracket OT \rrbracket.$$

Algebraic Semantics of MOF Metamodels. The $\mathbb{A}(\mathcal{M})$ theory constitutes the metamodel realization of the metamodel \mathcal{M} . This theory provides the model type $\llbracket \mathcal{M} \rrbracket$, which is syntactically represented by the sort Model in $\mathbb{A}(\mathcal{M})$. $\llbracket \mathcal{M} \rrbracket$ is the model type of collections of typed objects that have both a graph and a term structure. The semantics of the $\llbracket \mathcal{M} \rrbracket$ model type is defined by the equation

$$\llbracket \mathcal{M} \rrbracket = T_{\mathbb{A}(\mathcal{M}), \text{Model}}$$

and the *structural conformance relation* between a model definition M and its corresponding model type \mathcal{M} is then formally defined by the equivalence

$$M : \mathcal{M} \Leftrightarrow M \in \llbracket \mathcal{M} \rrbracket.$$

Embedding MOF Reflection into MEL Logical Reflection. We are now ready to discuss the third, final bootstrapping step, mentioned in Section 4.2. The logical reflective features of MEL [22,24], particularly its universal theory \mathcal{U} , make it possible to *internalize* the representation function mapping the specifications

of some other formalism to MEL as an equationally defined function *within* MEL (see [23]). In particular this can be done for our representation function for MOF

$$\llbracket \mathbb{A} \rrbracket : \text{MOF} \longrightarrow \text{SpecMEL}.$$

Note that the domain of the function \mathbb{A} , since it is the algebraic data type $\llbracket \text{MOF} \rrbracket = T_{\mathbb{A}(\text{MOF}), \text{Model}}$ defined by initial algebra semantics in Section 4.3, is already internalized within MEL. However, the set of (finitary) MEL specifications SpecMEL is a metalevel entity and therefore *outside* the object level of MEL, that is, not *directly* definable as an algebraic data type in MEL. However, this set is *representable* at the object level inside the universal theory \mathcal{U} by means of a faithful representation mapping

$$\text{SpecMEL} \longrightarrow \overline{\text{SpecMEL}} : (\Sigma, E) \mapsto (\overline{\Sigma}, \overline{E})$$

where, by definition, the set $\overline{\text{SpecMEL}}$ is the algebraic data type

$$\overline{\text{SpecMEL}} = T_{\mathcal{U}, \text{Module}}$$

corresponding to the elements of sort *Module* in the initial algebra of the universal theory \mathcal{U} . Then, the reflective internalization of the MOF algebraic semantics \mathbb{A} becomes an equationally defined function

$$\overline{\mathbb{A}} : \llbracket \text{MOF} \rrbracket \longrightarrow \overline{\text{SpecMEL}} : \mathcal{M} \mapsto \overline{\mathbb{A}(\mathcal{M})}$$

so that the term $\overline{\mathbb{A}(\mathcal{M})}$, for a specific metamodel \mathcal{M} , metarepresents the MEL theory $\mathbb{A}(\mathcal{M})$.

But since the universal theory \mathcal{U} faithfully simulates all the deductive capabilities of an object level theory (Σ, E) by means of its meta-representation $(\overline{\Sigma}, \overline{E})$ as data [24], this means that, given a metamodel \mathcal{M} in MOF, anything we can do with its associated MEL theory $\mathbb{A}(\mathcal{M})$ we can likewise perform at the metalevel with its meta-representation as data $\overline{\mathbb{A}(\mathcal{M})}$. Furthermore, in an efficient implementation of the relevant deductive functionality of \mathcal{U} such as the one provided in Maude's **META-LEVEL** module [22], the difference in computation time between performing a deduction in $\mathbb{A}(\mathcal{M})$ or performing the analogous deduction at the metalevel with $\overline{\mathbb{A}(\mathcal{M})}$ is negligible in practice.

In this way, the term $\overline{\mathbb{A}(\mathcal{M})}$, that for a specific metamodel \mathcal{M} meta-represents the MEL theory $\mathbb{A}(\mathcal{M})$, can then be used in different model-driven development scenarios, as illustrated in Section 5. The function \mathbb{A} is completely defined in [11] and is implemented as $\overline{\mathbb{A}}$ in MOMENT2. That the function $\overline{\mathbb{A}}$ is equationally definable follows from the general principle that \mathbb{A} is a computable function and that all computable functions can be equationally defined by a finite set of confluent and terminating equations [6]. In practical terms, this is just a matter of representing all the recursive definitions given for \mathbb{A} earlier in Section 4.4 as corresponding equational definitions at the metalevel.

4.5 MOMENT2 and the Eclipse Modeling Framework

MOMENT2 is a suite of tools, built on top of the Eclipse Modeling Framework (EMF), that provides support for formal model-driven development. The mathematical foundations of this tool suite directly rely on the algebraic semantics \mathbb{A} of MOF metamodels, presented in this work, which can be neatly exploited in the EMF, a widely used MOF-like metamodeling framework, by means of OMG standards, such as MOF, OCL and QVT. In this way, the formal analysis techniques that are presented in Section 5 can be applied to model-driven development practices in industrial environments, such as Rational Software Architect [30] or OSATE AADL [51]. In particular, MOMENT2 provides support for verifying metamodel conformance with OCL constraint satisfaction [16], QVT-like model transformations and their verification based on Maude’s reachability analysis and Maude’s LTL model checker, as explained in [15].

The integration of MOMENT2 into the EMF is based on a metamodel-independent, automatic bridge⁷ between EMF and Maude, so that EMF models M , such that $M : \mathcal{M}$, can be isomorphically represented as terms of sort *Model* in the theory $\mathbb{A}(\mathcal{M})$. That is, there is a mapping \mathbb{T} that formalizes an EMF model M that conforms to a metamodel \mathcal{M} as a term of sort *Model* in the theory $\mathbb{A}(\mathcal{M})$, and there is an inverse mapping \mathbb{T}^{-1} that parses a term of sort *Model* in the theory $\mathbb{A}(\mathcal{M})$ and provides its EMF representation back again so that $\mathbb{T}^{-1}(\mathbb{T}(M)) = M$. Since MOF metamodels \mathcal{M} are given as models $\mathcal{M} : \text{MOF}$ in EMF, \mathbb{T} can likewise be applied to them, so that the algebraic semantics \mathbb{A} of a MOF metamodel \mathcal{M} in the EMF is given by the theory $\mathbb{A}(\mathbb{T}(\mathcal{M}))$.

Fig. 3 shows a summary of the formalization of the metamodel \mathcal{M} of the example and a configuration of components as a model M in the formal framework, i.e., as a term of sort *Model*. In this way, MOMENT2 enables the prototyping and experimentation of formal model-driven development techniques and the formal analysis of model-based languages and models in Maude, as shown in Section 5.

5 Applications in Model-Driven Development Scenarios

In this Section we present some model-driven scenarios in which the algebraic semantics \mathbb{A} of MOF metamodels plays an important role. We illustrate the application of \mathbb{A} by using Maude and its support for formal reasoning in the running example of model-based software architecture reconfigurations. These techniques are internally used in MOMENT2, so that they can be applied in EMF-based environments through OMG standards and without explicit contact with the underlying MEL and Maude.

Section 5.1 presents some scenarios where \mathbb{A} provides the formal foundations for techniques and practices widely used in the model-driven community: metamodel conformance, semantics of MOF domain-specific languages, model

⁷ This bridge uses the API of the Maude Development Tools [37] for integrating Maude into Eclipse.

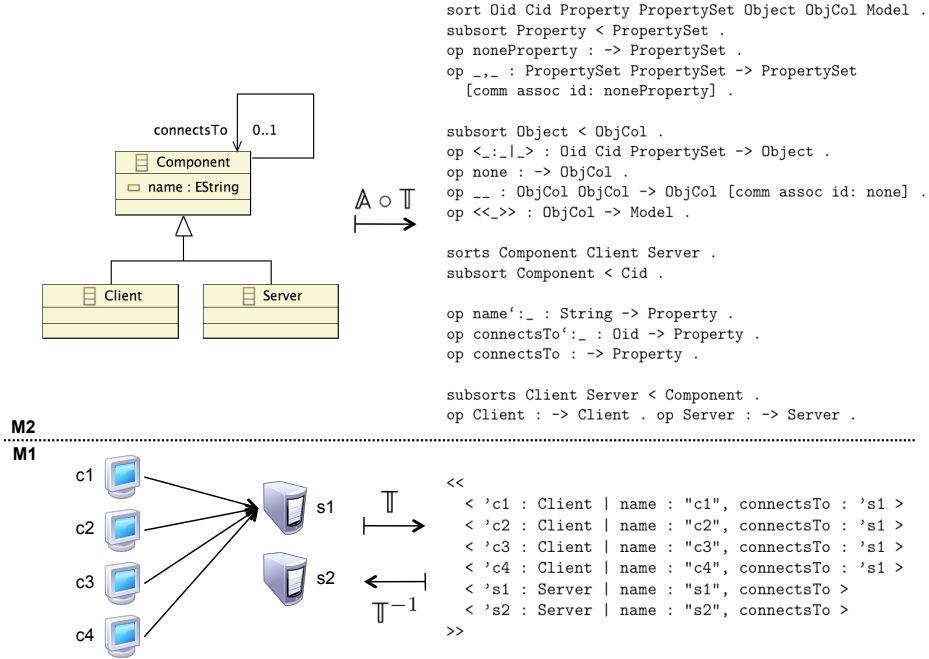


Fig. 3. EMF-Maude mappings in MOMENT2.

transformations and model management. Section 5.2 focuses on the formal implications and advantages of our approach where formal analysis is made available within modeling environments: static analysis based on structural constraint satisfaction, and dynamic analysis based on reachability analysis and LTL model checking.

5.1 Model-Driven Development Scenarios

Model-driven development techniques can be formalized in our approach as follows.

Model conformance. MOMENT2 provides an automated, deductive mechanism to check the structural conformance relation $M : \mathcal{M}$ by relying on Maude's implementation of MEL. Consider the metamodel \mathcal{M} and the model `conf1` in Fig. 3, given as a constant `conf1`. We can automatically check whether the model `conf1` conforms to the metamodel \mathcal{M} , formalized as the model type $[[\mathcal{M}]]$ and syntactically represented by the sort *Model* in the theory $\mathbb{A}(\mathcal{M})$, by evaluating the following membership in Maude:

```

red conf1 :: Model .
result Bool: true
    
```

Rewriting logic semantics of domain-specific languages. A MOF-based domain-specific language (DSL) provides modeling primitives that can be used to define concepts in a given domain. For example, real-time embedded systems can be specified with AADL [51], web service workflows can be defined with BPEL [44], or ontologies can be defined with OWL [35]. The abstract syntax of these languages is provided as MOF metamodels, their concrete syntax can be either textual or graphical, and their dynamic semantics can be defined either with approaches for defining the semantics of programming languages, with frameworks that provide library support for manipulating models, such as Kermeta [57], or with graph transformation systems [50,27,26]. The last one constitutes an interesting candidate for defining model-based DSLs due to the graph-based nature of models and to the formalization of the approach.

A MOF metamodel \mathcal{M} is used to define concepts in a particular domain, such as components in our running example, where software architectures are configurations of components that may be connected to each other through the *connectsTo* association end. In this section, we enrich the algebraic semantics $\mathcal{A}(\mathcal{M})$ with a rewrite rule to define a dynamic semantics for component configurations. In particular, we add a dynamic connection load balancing strategy, so that a server component should not have more than two connections at a time, i.e., when a component has more than two connections, the spare connections are forwarded to other components with less than two incoming connections. We depict the reconfiguration as a graph transformation rule in Fig. 4, where a rule is defined with a left-hand side (LHS) pattern and a right-hand side (RHS) pattern. Each pattern is constituted by nodes that represent *Component* objects in a model (graph) M and edges representing *connectsTo* references between them. A reconfiguration can be applied whenever the LHS pattern of the rule can be matched against a specific configuration M of components, and then the edges are manipulated as follows: an edge in the LHS and not in the RHS is removed from the configuration, an edge not in the LHS but in the RHS is added, the rest of edges remain unmodified. Marked edges indicate that the edges must not exist in order to apply the rule; this notion is known as a *negative application condition* in the graph transformation community.




Fig. 4. Reconfiguration rule.

The graph-theoretic nature of models is axiomatized in our algebraic semantics as a set of objects modulo the associativity, commutativity, and identity axioms of set union. The semantics of a reconfiguration can then be naturally

expressed as a *rewrite theory* [38] extending the algebraic semantics $A(\mathcal{M})$ of our metamodel specification with *rewrite rules* that are applied *modulo* the equational axioms. In this way, the above graph-transformation rule can be summarized at a high level as follows:

```

op free-reconfiguration : Model -> Model .
crl free-reconfiguration(M) =>
  free-reconfiguration(<< < O1 : C1 | PS1 >
    < O2 : C2 | connectsTo : O1, PS2 >
    < O3 : C3 | connectsTo : O1, PS3 >
    < O4 : C4 | connectsTo : O5, PS4 >
    < O5 : C5 | PS5 > OC >>)
if << < O1 : C1 | PS1 >
  < O2 : C2 | connectsTo : O1, PS2 >
  < O3 : C3 | connectsTo : O1, PS3 >
  < O4 : C4 | connectsTo : O1, PS4 >
  < O5 : C5 | PS5 > OC >> := M /\ nac(O5, M) .

```

where the expression $P := M$ matches the pattern P to the model variable M , $O1, O2, O3, O4, O5$: Oid , $C1, C2, C3, C4, C5$: $Component$, $PS1, PS2, PS3, PS4, PS5$: $PropertySet$, OC : $ObjCol$, M : $Model$, the LHS of the graph transformation rule in Fig. 4 corresponds to the pattern in the $P := M$ expression, its RHS corresponds to the RHS of the rewrite (after the $=>$ symbol), and the $nac(O5, M)$ condition corresponds to the negative application condition that enables the application of the rule:

```

op nac : Oid Model -> Bool .
eq [satisfiedNAC] : nac(O1, M) = true [owise] .
eq [counterexampleNAC] : nac(O1,
  << < O1 : C1 | PS1 >
    < O2 : C2 | connectsTo : O1, PS2 >
    < O3 : C3 | connectsTo : O1, PS3 > OC
  >>) = false .

```

Model transformations and operators for model management. Due to the executability of MEL specifications in Maude, the realization of MOF metamodels as MEL theories enhances the formalization and prototyping of model-driven development processes, such as:

- Model transformations [53], where translations of models between different modeling languages can be performed.
- Model-driven roundtrip engineering [10,46], where a model that constitutes the abstract syntax tree of a program can be translated into a model that specifies the software at a higher level of abstraction. Code may be generated again from recovered models in an automated way. An important issue is to keep both the abstract description and the code synchronized to deal with changes in a consistent way.

- Model traceability, where traceability is important to keep track of the changes performed by means of model transformations.
- Model management [7], where models can be manipulated by means of generic operators that rely on mappings between models. These operators permit, for example, merging models, generating mappings between models, and computing differences between models; they can be used to solve complex scenarios such as the roundtrip problem.

We illustrate the application of MOMENT2 to trace component reconfigurations, which can be applied to adapt a configuration to a given criteria, as the load balance strategy discussed above. We consider a traceability metamodel \mathcal{T} to define mappings between components as models T , such that $T : \mathcal{T}$. An equationally defined operator is presented to generate such models.

The metamodel \mathcal{T} in Fig. 5 permits defining mappings between components that may belong to different configurations as *TMapping* instances. A *TModel* instance constitutes the container for all the mappings between two models. Given configurations *conf1* and *conf2*, a trace model can be defined between them in order to map the elements that represent the same component in both configurations. Such mappings can be used to keep track of the changes that have been applied to the initial configuration or to identify structurally-equivalent elements in a model management scenario as presented in [12,14].

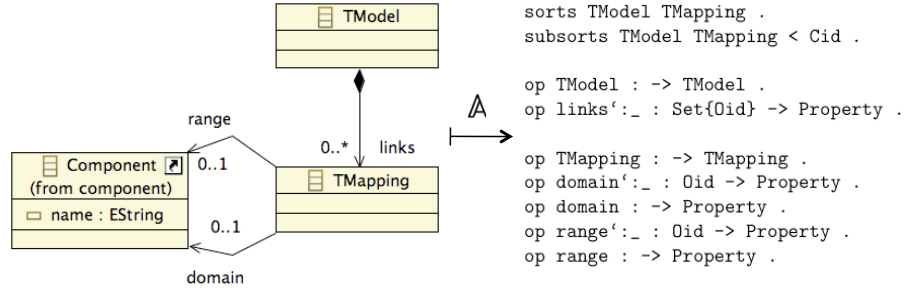


Fig. 5. Traceability metamodel \mathcal{T} and its algebraic realization $\mathbb{A}(\mathcal{T})$.

A *model transformation*, defined as an equationally defined function, can be used to automatically generate a trace model between two configurations. The following *match* operator takes two component configurations (the models *conf1* and *conf2* in Fig. 1) as inputs and generates a trace model with mappings between the components of both configurations. The *match* operator uses an auxiliary operator *\$match* that, in addition, takes a traceability model that is initialized with a *TModel* instance, and a natural number that is used to create object identifiers. The *\$match* operator adds a new link whenever two components with the same name are found. More complex structural criteria can be

used to match objects by simply traversing the objects that constitute the model by means of matching modulo associativity and commutativity axioms.

```

op match : Model Model -> Model{Trac} .
eq match(M1,M2) = $match(M1,M2,init-trac,0) .

op $match : Model Model Model{Trac} Nat -> Model{Trac} .
eq $match(M1,M2,TM, N) = TM [owise] .
eq $match( << < O1 : C1 | name : Name, PS1 > OC1 >>,
  << < O2 : C2 | name : Name, PS2 > OC2 >>,
  << < O3 : TModel | links : Set, PS3 > OC3 >>, N
) =
  $match( << OC1 >>, << OC2 >>,
    << < O3 : TModel | links : Set -> including( getNewOid(N) ), PS3 >
      < getNewOid(N) : TMapping | domain : O1, range : O2 > OC3 >>,
      N + 1 ) .
    
```

where $\text{Model}\{\text{Trac}\}$ is the sort of the model type $\llbracket \mathcal{S} \rrbracket$ in the theory $\mathbb{A}(\mathcal{S})$, init-trac is a constant that represents a trace model with a single $TModel$ instance, the getNewOid function obtains a new object identifier from a natural number, and the including operator adds an element to a set, as in OCL. The trace model between the configurations conf1 and conf2 can be generated by means of the command $\text{red match}(\text{conf1}, \text{conf2})$. The output model is illustrated in Fig. 6.

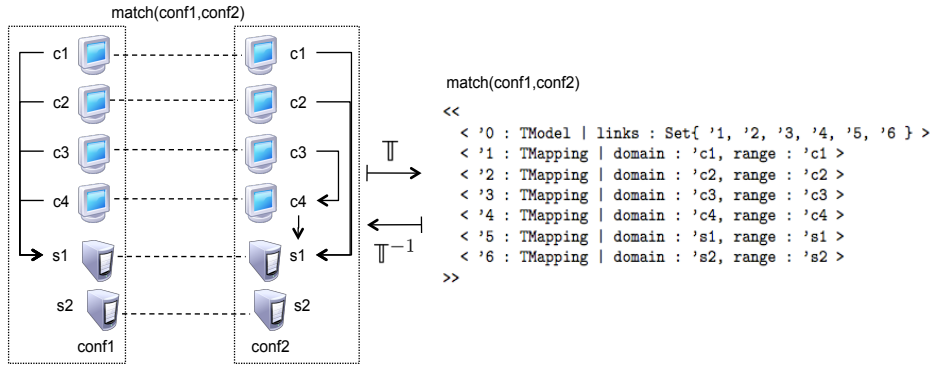


Fig. 6. Trace model and its term representation.

5.2 Formal Analysis Scenarios

In this subsection, we present some applications of the algebraic semantics \mathbb{A} for MOF metamodels to formal reasoning of model-driven development practices. In particular, static and dynamic analysis of model-based domain-specific languages (DSLs).

Static analysis. Once a metamodel \mathcal{M} is realized as a theory $\mathbb{A}(\mathcal{M})$, more axioms can be added to $\mathbb{A}(\mathcal{M})$ in order to enrich the semantics of the metamodel \mathcal{M} or to add semantics to the corresponding DSL. These axioms may correspond to OCL constraints as shown in [16] and can be used to verify that a specific model satisfies certain semantic properties.

For example, we can define a boolean predicate over a model that specifies that client components do not receive connections from other components in a specific configuration. This predicate is given as an equationally defined function that negates the satisfaction of the predicate when the property is not satisfied and that asserts its satisfaction in any other case:

```
op checkClientServer : Model -> Bool .
eq checkClientServer( << < O1 : C1 | connectsTo : O2, PS1 >
  < O2 : Client | PS2 > OC >> ) = false .
eq checkClientServer( M:Model ) = true [otherwise] .
```

where $O1, O2$: `Oid`, $C1$: `Component`, $PS1, PS2$: `PropertySet`, and OC : `ObjCol`. The models `conf1` and `conf2` in Fig. 1 can then be verified as follows:

```
red checkClientServer(conf1) .
result Bool: true

red checkClientServer(conf2) .
result Bool: false
```

where `conf2` does not satisfy the constraint.

Formal verification of behavioral specifications. The rewriting logic semantics of DSLs, defined by means of \mathbb{A} and additional rewrite rules, can also be used for formal dynamic analysis. Given a specific initial configuration M of components, we can use Maude's *search* command to model check whether or not all possible reconfigurations of an initial configuration preserve the constraint `checkClientServer`, introduced above.

The metamodel realization $\mathbb{A}(\mathcal{M})$, corresponding to Fig. 1, and the rewriting rule `free-reconfiguration` presented above define a state transition system, where states represent configurations M and M' of components and transitions $M \longrightarrow M'$ are given by one application of the reconfiguration rule. However, a reconfiguration of this kind could conceivably produce configurations M' of components that do not satisfy the constraint `checkClientServer`.

In Maude, the *search* command allows one to exhaustively explore (following a breadth-first strategy) the reachable state space defined by a state transition system as the one above, checking if an invariant is violated. We can use the *search* command to find out if the reconfiguration `free-reconfiguration` produces such an illegal configuration as follows⁸:

⁸ We have removed the types of the variables in the RHS of the command (after the `=>+` symbol) for the sake of simplicity.


```

search [1] free-reconfiguration(conf1) =>+
  free-reconfiguration(<< < 01 : C1 | connectsTo : 02, PS1 >
    < 02 : Client | PS2 > OC >>) .

```

where `conf1` is a constant that represents the initial configuration M in Fig. 1. This command finds a counterexample, `conf2` in Fig. 1, where a client component is connected to another client component. An alternative reconfiguration rule can be defined to avoid this problem as shown below, by indicating that the node `05` in the graph patterns of the rule in Fig. 4 is of type *Server*. No counterexamples are found when running again the search command with the new reconfiguration rule.

```

op safe-reconfiguration : Model -> Model .
crl safe-reconfiguration(M) =>
  safe-reconfiguration(<< < 01 : C1 | PS1 >
    < 02 : C2 | connectsTo : 01, PS2 >
    < 03 : C3 | connectsTo : 01, PS3 >
    < 04 : C4 | connectsTo : 05, PS4 >
    < 05 : Server | PS5 > OC >>)
if << < 01 : C1 | PS1 >
  < 02 : C2 | connectsTo : 01, PS2 >
  < 03 : C3 | connectsTo : 01, PS3 >
  < 04 : C4 | connectsTo : 01, PS4 >
  < 05 : Server | PS5 > OC >> := M /\ nac(05, M) .

```

Linear temporal logic model checking. Maude also provides a model checker where properties can be given as Linear Temporal Logic (LTL) formulae [22]. Taking into account the `safe-reconfiguration` rule that has been added to the theory $\mathcal{A}(\mathcal{M})$, we can also model check liveness properties, such as the fact that the server $s1$ will always have a balanced load L eventually. This property can be formulated in LTL as $\square \diamond \text{balanced}("s1", L)$, where $\text{balanced}("s1", L)$ is a state predicate that is satisfied when the server component with name "s1" has L or less connections. Following the guidelines provided in [22], we defined a subtheory inclusion to use Maude's model checker, defined in the theory `MODEL-CHECKER`, into the theory $\mathcal{A}(\mathcal{M})$ that is extended with the rule `safe-reconfiguration`. The sort `Model` is defined as subsort of the sort `State`, used by the model checker to represent system states. The predicate symbol balanced is then defined as a parametric predicate of sort `Prop` so that the satisfaction relation $_|_$ between system states and state predicates can be equationally defined for the case in which the state predicate is satisfied:

```

subsort Model < State .

op balanced : String Nat -> Prop .
eq free-reconfiguration( << < 01 : Server | name : Name, PS1 > OC >> )
  |= balanced(Name, L)
= (countBalance(01, << < 01 : Server | name : Name, PS1 > OC >>) <= L) .

```

where `countBalance` is an equationally defined function that counts the number of connections to a given component with identifier `O1` in the model `M`:

```
op countBalance : Oid Model -> Nat .
eq countBalance( O1, M ) = 0 [owise] .
eq countBalance( O1, << < O2 : C2 | connectsTo : O1, PS2 > OC >> ) =
  1 + countBalance(O1, << OC >> ) .
```

The model predicate `balanced("s1",2)` can be used in a LTL formula to model check that the server "s1" will end up with 2 or less connections as follows:

```
red modelCheck(safe-reconfiguration(conf1), [] <> balanced("s1", 2)) .
```

6 Related work

The meaning of the *metamodel* notion has been widely discussed in the literature, see for example [36,52,34,48]. There is a consensus that a metamodel can play several roles: as *data*, as *type* or as *theory*. In this paper, we have formally expressed each of these roles by means of the notions of metamodel \mathcal{M} , model type $\llbracket \mathcal{M} \rrbracket$, and metamodel realization $\mathbb{A}(\mathcal{M})$, respectively.

The current MOF standard does not provide any guidelines to implement a reflective mechanism that obtains the semantics of a metamodel. An informal attempt to realize MOF metamodels as Java programs is provided in the Java Metadata Interface (JMI) specification [32], which is defined for a previous version of the MOF standard. A mapping of this kind has been successfully implemented in modeling environments such as the Eclipse Modeling Framework. By contrast, our \mathbb{A} function gives us an executable formal specification of the algebraic semantics of any EMOF metamodel \mathcal{M} .

Although EMOF metamodels can be viewed as simplified UML class diagrams, formal approaches for metamodeling need a reflective mechanism, such as \mathbb{A} in our approach, to provide the semantics of modeling languages. This mechanism is not needed in UML, where the modeling language is fixed. We focus on several approaches for metamodeling that rely on different formalisms.

The Meta-Modeling Language is a meta-circular language based on the MML calculus [20], which provides an operational semantics for both UML modeling constructs and OCL operators. Modeling languages can be precisely defined in MML by explicitly specifying its abstract syntax, its *semantic domain* and a mapping between the concepts involved in both [21]. This mapping can be viewed as the application of the function \mathbb{A} to a specific metamodel \mathcal{M} .

Alloy [31] is a declarative language based on first-order relational logic in which systems with constraints can be modeled. The Alloy analyzer [1] provides an automated mechanism for constraint satisfaction with two main functionalities: *simulation* for producing valid instances of an Alloy specification and *assertion* for verifying constraints. [4] provides an encoding of EMOF metamodels with OCL constraints into Alloy so that the Alloy analyzer [1] is used to generate models that conform to a metamodel (automated test case generation)

and to verify that OCL constraints can be satisfied. Counterexamples and logical inconsistencies are found when the constraints are not satisfied. In [5], the Alloy analyzer is used to verify relational model transformations in order to ensure that a model transformation cannot produce invalid models. However, Alloy has a simple type system where only integers can be used in attribute values.

In [47], constructive type theory is used for defining a typed metamodeling framework, where models, which are defined as terms, can also be represented as types by means of a reflection mechanism such as \mathbb{A} . In this framework, the conformance relation is implicitly provided by construction: only valid models can be defined as terms, and their definition constitutes a formal proof of the fact that the model belongs to the corresponding type by means of the Curry-Howard isomorphism.

In the graph transformations field [50,27,26], metamodels are defined as type graphs with node inheritance, and models are defined as attributed typed graphs. The main difference between type graphs and metamodels rely on the use of composition associations in EMOF metamodels, which can be used for defining hierarchies of composite objects in models. A notion of *graph with containments* is introduced in [9], where the authors show how graph transformations can be used as a formal backend for model transformations. In this way, the theory on graph transformations and related tools can be used to perform formal analysis of model transformations. In particular, the authors show how to analyse termination and confluence of model transformations that are encoded as graph transformations in the algebraic graph transformation environment AGG [2].

There are a number of metamodeling approaches based on Maude. Maude already provides support for object-oriented programming [39], where objects, the *isInstanceOf* relation and the *class specialization* relation are supported. The dynamics of object-oriented systems can be provided by means of term rewriting.

The static semantics of the UML metamodel (version 1.3) has been previously formalized as an algebraic specification in MEL [29]. In this approach, the authors already took the MOF approach into account, although the MOF standard was in its early stages. In [28,3], the authors provide a formal four-layered framework where: (i) some parts of the MOF meta-metamodel are formalized in a MEL theory at M3 level (called MOF layer); (ii) the UML class diagram and the object diagram metamodels are provided as MEL theories, called *syntactic specification* and *semantic specification* respectively, at M2 level (called UML metamodel layer); (iii) UML class diagrams are defined as terms in the syntactic specification theory at M1 level (called *domain model* layer); and (iv) object diagrams are defined as terms in the semantic specification theory at M0-level (named *user objects* layer). A novel feature in this approach relies on the reuse of the reflective facilities of MEL to provide support for the evolution of UML-based software artifacts [56]. The authors focused on the verification of static properties by using Maude as an implementation of MEL and the language to define the constraints.

Another approach [49] based on Maude uses the KM3 language [33] for indirectly defining EMOF metamodels. In this work, the authors present how KM3

specifications of metamodels can be represented as object modules [22] in Full-Maude and how models can be defined as Maude collections of objects. However, no automated support is provided for representing models as terms since KM3 only permits defining the textual concrete syntax of metamodels. That is, mappings like \mathbb{T} and \mathbb{T}^{-1} are not defined for models that conform to metamodels that are extracted from KM3 specifications. In this way, the user has to define the models in Maude notation directly. The authors provide a mechanism to represent KM3 specifications of metamodels as collections of objects at a syntactic level so that Maude is used to statically analyse KM3 specifications: to check when two metamodels describe model subtypes, to infer metamodels from models and to compute metrics. However, model types are not algebraically characterized.

Our algebraic semantics \mathbb{A} for EMOF metamodels \mathcal{M} in MEL formally defines the notions of *metamodel realization* $\mathbb{A}(\mathcal{M})$, *model type* $\llbracket \mathcal{M} \rrbracket$ and *model conformance* $M : \mathcal{M}$. This means that \mathbb{A} enables reasoning with model types at an algebraic level and not just at a syntactical level. Due to the graph nature of models, the algebraic semantics \mathbb{A} for MOF can also be used as an algebraic environment for graph transformations, where Maude’s analysis capabilities, such as reachability analysis and LTL model checking, can be reused. The complete algebraic formalization of EMOF metamodels together with OCL can be found in [11], where composition associations are also taken into account in the formalization. The algebraic semantics \mathbb{A} is implemented in MOMENT2 where EMF is used as implementation of the EMOF standard. Furthermore, the generic mappings \mathbb{T} and \mathbb{T}^{-1} allow representing EMF models as terms in MOMENT2 in a transparent way to the user. This is an essential feature in MOMENT2, where the goal is to apply Maude for formal model management tasks by using OMG standards, such as MOF, OCL and QVT.

7 Conclusions and Future Work

In this work we have proposed an algebraic semantics for the MOF metamodeling framework, formalizing notions not yet clear in the MOF standard. In our approach, we give an explicit formal representation for each of the different notions that may be involved in a metamodeling framework: model type $\llbracket \mathcal{M} \rrbracket$, metamodel realization $\mathbb{A}(\mathcal{M})$, and metamodel conformance $M : \mathcal{M}$. Our work provides an algebraic executable formalization of the MOF standard that can be reused in standard-compliant frameworks.

This algebraic framework opens a wide spectrum of interesting applications for model-driven development. In particular, we have shown how it can be used for automatically checking metamodel conformance, defining domain-specific languages and specifying model transformations and model management operators. In addition, Maude’s formal verification facilities can be used for static and dynamic analysis of domain-specific languages, such as checking constraints over models, reachability analysis and LTL model checking.

The algebraic semantics for MOF provides the foundational notions for a model management tool suite, MOMENT2 [43], that supports both OCL [16] and QVT-like model transformations [15]. MOMENT2 uses MOF, OCL and a QVT-like model transformation language as interface, so that the techniques that have been illustrated in this work are internally used in a transparent way to the user. For the development of this framework, we have relied on the experience gained in previous prototypes that gave algebraic executable specifications for OCL [17], QVT [13] and model management operators [12,14]. In addition, grammar-based software artifacts can also be related to models by specifying context-free grammars as MEL signatures. This last feature makes our framework also suitable for forward and reverse Model-Driven Engineering.

In future work, we plan to apply the algebraic MOF framework together with the aforementioned tool for model transformations and Maude-based formal reasoning techniques in model-driven development scenarios, where software systems that are developed contain critical properties that have to be verified. In particular, we are considering the formal analysis of real-time embedded systems in the avionics domain, by using model-based languages like the Architecture Analysis and Design Language (AADL) [51].

Acknowledgments. This work has been partially supported by the project META TIN2006-15175-C05-01, by the ONR Grant N00014-02-1-0715, by NSF Grant IIS-07-20482, and by the project SENSORIA, IST-2005-016004.

References

1. Alloy Analyzer Homepage, 2008. <http://alloy.mit.edu/>.
2. AGG Homepage. <http://tfs.cs.tu-berlin.de/agg/>.
3. J. L. F. Alemán. *A formalization proposal of the UML four-layered architecture*. PhD thesis, Murcia University, April 2002. (In Spanish).
4. K. Anastasakis, B. Bordbar, G. Georg, and I. Ray. Uml2alloy: A challenging model transformation. In G. Engels, B. Opdyke, D. C. Schmidt, and F. Weil, editors, *MoDELS*, volume 4735 of *LNCS*, pages 436–450. Springer, 2007.
5. K. Anastasakis, B. Bordbar, and J. M. Küster. Analysis of Model Transformations via Alloy. In *4th International Workshop on Model Driven Engineering, Verification, and Validation*, volume 5002 of *LNCS*, pages 145–150. Springer, 2008.
6. J. Bergstra and J. Tucker. Characterization of computable data types by means of a finite equational specification method. In J. W. de Bakker and J. van Leeuwen, editors, *Automata, Languages and Programming, Seventh Colloquium*, pages 76–90. Springer-Verlag, 1980. *LNCS*, Volume 81.
7. P. A. Bernstein. Applying Model Management to Classical Meta Data Problems. In *Proceedings of the 1st Biennial Conference on Innovative Data Systems Research (CIDR)*, 2003.
8. J. Bézivin. On the unification power of models. *Software and System Modeling (SoSym)*, 4(2):171–188, 2005.
9. E. Biermann, C. Ermel, and G. Taentzer. Precise Semantics of EMF Model Transformations by Graph Transformation. In K. Czarnecki, editor, *MoDELS'08*, *LNCS*. Springer, 2008. To appear.

10. M. Bork, L. Geiger, C. Schneider, and A. Zündorf. Towards roundtrip engineering - a template-based reverse engineering approach. In I. Schieferdecker and A. Hartman, editors, *ECMDA-FA*, volume 5095 of *LNCS*, pages 33–47. Springer, 2008.
11. A. Boronat. *MOMENT: a formal framework for MOdel manageMENT*. PhD in Computer Science, Universitat Politècnica de València (UPV), Spain, 2007. http://www.cs.le.ac.uk/people/aboronat/papers/2007_thesis_ArturBoronat.pdf.
12. A. Boronat, J. A. Carsí, and I. Ramos. Automatic Support for Traceability in a Generic Model Management Framework. In A. Hartman and D. Kreische, editors, *ECMDA-FA*, volume 3748 of *LNCS*, pages 316–330. Springer, 2005.
13. A. Boronat, J. A. Carsí, and I. Ramos. Algebraic specification of a model transformation engine. In L. Baresi and R. Heckel, editors, *FASE*, volume 3922 of *LNCS*, pages 262–277. Springer, 2006.
14. A. Boronat, J. A. Carsí, I. Ramos, and P. Letelier. Formal model merging applied to class diagram integration. *Electr. Notes Theor. Comput. Sci.*, 166:5–26, 2007.
15. A. Boronat, R. Heckel, and J. Meseguer. Rewriting Logic Semantics and Verification of Model Transformations. Technical Report CS-08-004, University of Leicester, 2008. <http://www.cs.le.ac.uk/people/aboronat/papers/boHeMe-rl-mt.pdf>.
16. A. Boronat and J. Meseguer. Algebraic Semantics of OCL-constrained Metamodel Specifications. Technical Report UIUCDCS-R-2008-2995, UIUC, 2008. <http://www.cs.uiuc.edu/research/techreports.php?report=UIUCDCS-R-2008-2995>.
17. A. Boronat, J. Oriente, A. Gómez, I. Ramos, and J. A. Carsí. An Algebraic Specification of Generic OCL Queries Within the Eclipse Modeling Framework. In A. Rensink and J. Warmer, editors, *ECMDA-FA*, volume 4066 of *LNCS*, pages 316–330. Springer, 2006.
18. A. Bouhoula, J.-P. Jouannaud, and J. Meseguer. Specification and proof in membership equational logic. *Theoretical Computer Science*, 236:35–132, 2000.
19. I. Cervesato and M.-O. Stehr. Representing the MSR cryptoprotocol specification language in an extension of rewriting logic with dependent types. In P. Degano, editor, *Proc. Fifth International Workshop on Rewriting Logic and its Applications (WRLA'2004)*, volume 117. Elsevier ENTCS, 2004. Barcelona, Spain, March 27 - 28, 2004.
20. T. Clark, A. Evans, and S. Kent. The metamodelling language calculus: Foundation semantics for uml. In H. Hußmann, editor, *FASE*, volume 2029 of *LNCS*, pages 17–31. Springer, 2001.
21. T. Clark, A. Evans, and S. Kent. Engineering modelling languages: A precise metamodelling approach. In *FASE '02: Proceedings of the 5th International Conference on Fundamental Approaches to Software Engineering*, volume 2306 of *LNCS*, pages 159–173, London, UK, 2002. Springer.
22. M. Clavel, F. Durán, S. Eker, J. Meseguer, P. Lincoln, N. Martí-Oliet, and C. Talcott. *All About Maude*. Springer LNCS Vol. 4350, 2007.
23. M. Clavel, F. Durán, S. Eker, J. Meseguer, and M.-O. Stehr. Maude as a formal meta-tool. In J. Wing and J. Woodcock, editors, *FM'99 — Formal Methods*, volume 1709 of *Springer LNCS*, pages 1684–1703. Springer-Verlag, 1999.
24. M. Clavel, J. Meseguer, and M. Palomino. Reflection in membership equational logic, many-sorted equational logic, Horn logic with equality, and rewriting logic. *Theoretical Computer Science*, 373:70–91, 2007.
25. Eclipse Organization. The Eclipse Modeling Framework, 2007. <http://www.eclipse.org/emf/>.

26. H. Ehrig, K. Ehrig, U. Prange, and G. Taentzer. *Fundamentals of Algebraic Graph Transformation (Monographs in Theoretical Computer Science. An EATCS Series)*. Springer, March 2006.
27. H. Ehrig, U. Montanari, H.-J. Kreowski, G. Rozenberg, and H.-J. Kreowski. *Handbook of Graph Grammars and Computing by Graph Transformations, Vol. 3: Concurrency, Parallelism, and Distribution*. World Scientific Publishing Company, July 1999.
28. J. L. Fernández and A. Toval. Can intuition become rigorous? Foundations for UML model verification tools. In *International Symposium on Software Reliability Engineering (ISSRE 2000)*. IEEE, October 2000. San Jose, California, USA.
29. J. L. Fernández and A. Toval. *Seamless Formalizing the UML Semantics through Metamodels*. Unified Modeling Language: Systems Analysis, Design, and Development Issues. Idea Group Publishing, 2001.
30. IBM. Rational Software Architect for WebSphere Software, 2008. <http://www-01.ibm.com/software/awdtools/swarchitect/websphere/>.
31. D. Jackson. *Software Abstractions: Logic, Language, and Analysis*. The MIT Press, 2006.
32. Java Community Process. The Java Metadata Interface (JMI) Specification (JSR 40), 2002. <http://www.jcp.org/en/jsr/detail?id=40>.
33. F. Jouault and J. Bézivin. Km3: A dsl for metamodel specification. In R. Gorrieri and H. Wehrheim, editors, *FMOODS*, volume 4037 of *LNCS*, pages 171–185. Springer, 2006.
34. T. Kühne. Matters of (meta-) modeling. *Software and Systems Modeling (SoSyM)*, 5:369–385(17), December 2006.
35. L. W. Lacy. *Owl: Representing Information Using the Web Ontology Language*. Trafford Publishing, January 2005.
36. J. Ludewig. Models in software engineering - an introduction. *Inform., Forsch. Entwickl.*, 18(3-4):105–112, 2004.
37. Maude Development Tools, 2006. <http://moment.dsic.upv.es>.
38. J. Meseguer. Conditional rewriting logic as a unified model of concurrency. *Theoretical Computer Science*, 96(1):73–155, 1992.
39. J. Meseguer. A logical theory of concurrent objects and its realization in the Maude language. In G. Agha, P. Wegner, and A. Yonezawa, editors, *Research Directions in Concurrent Object-Oriented Programming*, pages 314–390. MIT Press, 1993.
40. J. Meseguer. Membership algebra as a logical framework for equational specification. In F. Parisi-Presicce, editor, *Proc. WADT'97*, pages 18–61. Springer LNCS 1376, 1998.
41. J. Meseguer and G. Roşu. Rewriting logic semantics: From language specifications to formal analysis tools. In *Proc. Intl. Joint Conf. on Automated Reasoning IJCAR'04, Cork, Ireland, July 2004*, pages 1–44. Springer LNAI 3097, 2004.
42. J. Meseguer and G. Roşu. The rewriting logic semantics project. *Theoretical Computer Science*, 373(3):213–237, 2007.
43. MOMENT2 Homepage, 2008. <http://www.cs.le.ac.uk/people/aboronat/tools/moment2>.
44. OASIS. Web Services Business Process Execution Language Version 2.0, 2007. <http://docs.oasis-open.org/wsbpel/2.0/wsbpel-v2.0.pdf>.
45. Object Management Group. Meta Object Facility (MOF) 2.0 Core Specification (ptc/06-01-01), 2006. <http://www.omg.org/cgi-bin/doc?formal/2006-01-01>.
46. E. V. Paesschen, W. D. Meuter, and M. D'Hondt. Selfsync: a dynamic round-trip engineering environment. In *OOPSLA '05: Companion to the 20th annual ACM*

- SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 146–147, New York, NY, USA, 2005. ACM.
47. I. Poernomo. The meta-object facility typed. In H. Haddad, editor, *SAC*, pages 1845–1849. ACM, 2006.
 48. A. Rensink. Subjects, models, languages, transformations. In J. Bézivin and R. Heckel, editors, *Language Engineering for Model-Driven Software Development*, volume 04101 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, 2004.
 49. J. R. Romero, J. E. Rivera, F. Durán, and A. Vallecillo. Formal and Tool Support for Model Driven Engineering with Maude. *Journal of Object Technology*, 6(9):187–207, 2007.
 50. G. Rozenberg. *Handbook of Grammars and Computing by Graph Transformation, Vol. 1: Foundations*. World Scientific Publishing Company, January 1997.
 51. SAE. AADL, 2007. <http://www.aadl.info/>.
 52. E. Seidewitz. What models mean. *Software, IEEE*, 20(5):26–32, 2003.
 53. S. Sendall and W. Kozaczynski. Model transformation: The heart and soul of model-driven software development. *IEEE Software*, 20(5):42–45, 2003.
 54. T. Serbanuta, G. Roşu, and J. Meseguer. A rewriting logic approach to operational semantics. To appear in *Information and Computation*, 2008.
 55. M.-O. Stehr and J. Meseguer. Pure type systems in rewriting logic: Specifying typed higher-order languages in a first-order logical framework. In *Essays in Memory of Ole-Johan Dahl*, pages 334–375. Springer LNCS Vol. 2635, 2004.
 56. A. Toval and J. L. Fernández. Formally modeling uml and its evolution: a holistic approach. *Kluwer Academic Publishers*, September 2000. FMOODS’00, Formal Methods for Open Object-Based Distributed Systems. Stanford, California, USA.
 57. Triskell Team. Kermeta, 2008. <http://www.kermeta.org/>.