ModHel'X: A Component-Oriented Approach to Multi-Formalism Modeling

Cécile Hardebolle[1]   Frédéric Boulanger[2]

Supélec – Départment Informatique
91192 Gif-Sur-Yvette cedex, FRANCE

Abstract

We present an approach to multi-formalism modeling which addresses two important issues in this field: to provide support for the specification of the semantics of a modeling formalism, and to allow the specification of the interactions between parts of a model described using different modeling formalisms. ModHel'X focuses on the execution of models, which we consider as the computation of one possible behavior of the model. This includes simulation, code generation or real-time execution. Based on the concept of Model of Computation, ModHel'X includes a UML meta-model which defines a component-oriented and hierarchical abstract syntax for describing the structural elements of a modeling language. The semantics of a modeling language, i.e. the corresponding model of computation, is expressed in an imperative style and addresses three aspects: control, data and time. The semantic adaptation between the heterogeneous levels found in the hierarchy of a model is described using the same imperative syntax. Finally, an execution engine has been developed for ModHel'X, which is able to interpret multi-formalism models for simulation purposes.

Keywords: Multi-formalism Modeling, Heterogeneous Modeling, Multi-Paradigm Modeling, Model of Computation, Meta-model, Simulation, Model Transformation

1   Introduction

Complex systems are inherently heterogeneous because of the diverse nature of their numerous parts: hardware, software, digital, analog, internal or external IPs (Intellectual Properties), etc. Modeling such systems requires multiple modeling formalisms, adapted to the nature of each part of the system, the aspect on which the model focuses (functionality, time, power consumption. . .) and to the level of abstraction at which the system, or one of its parts, is studied. As emphasized by [9], having a global model of such a system all along the design process is necessary in order to answer questions about properties of the whole system, and in particular about its behavior. Such a model is said to be *multi-formalism* [18].

Multi-formalism modeling (or heterogeneous modeling) is an emerging field and different aspects of it have been studied: mathematical foundations [14], tools for validation [2] or simulation [10]. A central problem is to establish the meaning of the composition of heterogeneous parts of a model and to ensure their correct interoperation when using the model to answer questions about the designed system [12].

We believe that the first step to be taken for solving this problem is to provide means for the precise specification of the semantics of the modeling formalisms that we want to use. Indeed, except for a few mathematically founded languages, the semantics of a modeling language is often described using natural language, what may lead to ambiguities and to diverse interpretations by different tools along the design chain. When combining different modeling languages in a model, ambiguities in the semantics of one of them make it impossible to define the overall semantics of the model. In this context, semantic variations as found in UML are acceptable only if the variation used is explicitly stated.

In ModHel'X, we propose a set of tools for allowing the precise specification of the semantics of a modeling formalism without referring to any model instance (i.e. at the meta-modeling level [18]). In order to facilitate the combination of multiple modeling languages in models, our approach relies on concepts of component-oriented and hierarchical modeling [4]. The encapsulation principle is a major advantage for heterogeneous modeling since its very purpose is to hide the internal mechanisms of the components. In this context, hierarchy is a structural way of combining the heterogeneous parts of a model, as well as a simple abstraction mechanism. In addition, we provide means for explicitly describing how to adapt the semantics between heterogeneous parts of a model in order to obtain a meaningful multi-formalism model of a system. We believe that the semantic adjustment between parts that use different modeling languages cannot be completely determined automatically. Even if usual interaction patterns between modeling languages often exist, they are not unique. ModHel'X permits the description of such patterns and allows the designer to choose the most suitable in a given model. These descriptions may be reused in different contexts and parameters allow their flexible adaptation.

Finally, we have developed for ModHel'X an execution engine which is able to interpret multi-formalism models and to simulate their behavior.

The remainder of the paper is organized as follows. In Section 2 we review some of the related work and motivate our approach. Section 3 details and illustrates the main principles of ModHel'X. We discuss some specific aspects of our approach in Section 4, before concluding.

2 Existing multi-formalism approaches and motivations

In meta-modeling approaches such as Kermeta [19], the abstract syntax of a modeling language is described as a UML meta-model. The elements of this meta-model have methods whose semantics is defined in an imperative language. Each modeling language has a different meta-model in Kermeta. In the context of heterogeneous modeling, the definition of the combination of several modeling languages using such approaches implies either the definition of a meta-model which is the union of all the meta-models of the involved languages or the definition of transforma-

tions from each meta-model to a meta-model chosen among them. Defining a union meta-model seems neither reasonable nor scalable since it implies modification of the meta-model and the associated model transformations when an additional modeling language is taken into consideration. The second method is much more interesting since it is more flexible: the target meta-model can be chosen according to the question to be answered about the system. Such an approach is implemented in the $ATOM^3$ tool [7]. However, the way the different heterogeneous parts of the model are "glued" together does not seem to be addressed by this approach.

Other approaches [15,13] are also based on model transformation. In particular, [13] states that it is possible to formally define the semantics of a modeling language by defining a mapping to an already existing formally defined modeling language.

Another approach for defining the semantics of a modeling language, is to define the constructs of the language in a fixed abstract syntax — or meta-model — which is component oriented (as in [4]) and to consider that the semantics of a modeling language is given by its "Model of Computation (MoC)". Such an approach is implemented in Ptolemy [9]. A model of computation (called "domain" in Ptolemy) is a set of rules for interpreting the relations between the components of a model. In this approach, the meta-model is the same for each language and what defines the semantics of the language is the way the elements of this meta-model are interpreted by the corresponding MoC. Heterogeneous models are organized into hierarchical layers, each one involving only one MoC. Thanks to this architecture, MoCs (i.e. modeling languages) are combined in pairs at the boundary between two hierarchical levels. The main drawback of the Ptolemy approach is that the way MoCs are combined at a boundary between two hierarchical levels is fixed and coded into the Ptolemy kernel. This implies that a modeler has either to rely on the default adaptation performed by the tool, or to modify the design of parts of its model (by adding adaptation components) in order to obtain the behavior he expects.

The approach we propose is based on the concept of *model of computation (MoC)* as defined in [9]. Our UML meta-model, which is inspired by the abstract syntax of Ptolemy, contains special constructs for making the interactions between heterogeneous MoCs explicit and easy to define. In order to interpret a model in ModHel'X, it is necessary to describe its structure using our meta-model. Then, we define an interpretation of the elements of our meta-model which matches the semantics of the original language. Such an interpretation is what we call a Model of Computation. The interpretation of a model according to a MoC gives the same behavior as the interpretation of the original model according to the semantics of its modeling language. The same concepts used to define MoCs are used to define how different MoCs are "glued" together in heterogeneous models, at the boundary between two hierarchical layers. The execution engine of ModHel'X relies on the precise specification of the models of computation and of their interactions to determine without ambiguity the behavior of multi-formalism models.

Other approaches of heterogeneous modeling are also based on a hierarchical and component oriented abstract syntax. BIP (Behavior, Interaction, Priority) [2] provides formally defined mechanisms for describing combinations of components in a model using heterogeneous interactions. BIP does not consider components

as black boxes and has access to the description of their behavior. This allows the formal verification of properties on the model. It is important to note that, in BIP, the description of the interactions between components is made at the M1 level.

The "42" approach [16] seems closer to ours. Based on the synchronous paradigm, 42 generates the code of the MoCs (called "controllers") from the contracts of the components (described using automata), the relations between their ports and additional information related to activation scheduling. The strength of this approach relies on the description of the behavioral contract of components. However, such a description may not be available (in the case of an external IP for instance) or may not be easy to establish, in the case of continuous time behaviors for example.

Metropolis [1] also relies on the concept of model of computation, but it focuses on MoCs related to process networks. It originates from trace algebras [5] and is closely related to the tag semantics approaches [14,3], which provide mathematical frameworks for the formalization of MoCs and their interactions but are very far from model execution.

3   Modeling heterogeneous systems with ModHel'X

3.1   Black boxes and snapshots

We adopt a component-oriented approach in which we consider components as black boxes, called *blocks*, in order to decouple the internal model of a component from the model of the system in which it is used. Therefore, the behavior of a block is observable only at its interface: nothing is known about what is happening inside the block, and in particular whether the block is even computing something.

In addition, instead of "triggering" the behavior of a block, we only *observe* its interface. When we need to observe a block, we ask it to provide us with a coherent view of its interface at this moment. A block can therefore be active even when we do not observe it. This is a key point in our approach because it allows us to embed asynchronous processes in a model without synchronizing them: we simply observe them at instants suitable for the embedding model. The behavior of a block or a model is therefore a sequence of observations. An observation of a model is defined as the combination of the observations of its blocks according to a MoC. This definition holds at all the levels of a hierarchical model. The observation of the top-level model, i.e. the model of the overall system, is a *snapshot* [6] which defines the exact state of the interface of each block at a given instant (such a notion is also defined in the context of UML [21]). We detail the way a snapshot is obtained using the rules expressed by a MoC in Section 3.4.

3.2   Time

The notions of time used in different models of computation are varied (real time, logical clocks, partial order on signal samples, etc.), and ModHel'X must support all of them. Moreover, in an heterogeneous model, different notions of time are combined and each part of the model may have its own "date" in a given snapshot. Therefore, the succession of snapshots is the only notion of time which is shared by
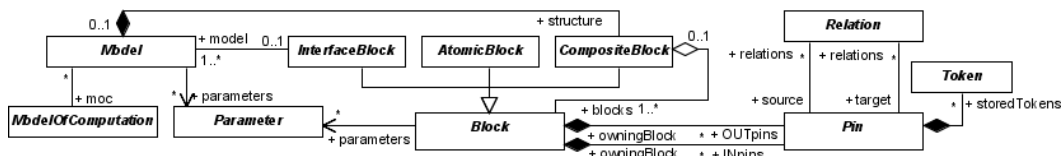
Fig. 1. Generic meta-model for representing the structure of models

all MoCs and which is predefined in ModHel'X. On this sequence of instants, each MoC can define its own notion of time.

A snapshot of a model is made whenever its environment (i.e. the input data) changes, but also as soon as any block at any level of the hierarchy needs to be observed because its state has changed. To this end, each component of an heterogeneous model can give constraints on its own date at the next snapshot. For instance, in a timed automaton, a time out transition leaving the current state must be fired even if no input is available. This can be achieved by requiring, when entering this state, that the next snapshot occurs before the timeout expires. This feature is a major departure from the Ptolemy approach, where the root model drives the execution of the other layers of the hierarchy.

Times in two MoCs may be synchronized by the interaction pattern at the boundary of two hierarchical levels. Thus, time constraints can propagate through the hierarchy up to the top level model.

3.3   A generic meta-model for representing the structure of models

The generic meta-model that we propose, shown on Figure 1, defines abstract concepts for representing the structural elements of models. Each of these concepts can be specialized in order to represent notions that are specific to a given modeling language, but their semantics is given by the MoCs which interpret them.

In the *structure* of a *model*, *blocks* are the basic units of behavior. *Pins* define the interface of models and blocks. The interactions between blocks are represented by *relations* between their pins. Relations are unidirectional and do not have any behavior: they are interpreted according to the MoC in order to determine how to combine the behaviors of the blocks they connect. For instance, a relation can represent a causal order between two blocks as well as a communication channel.

In Modhel'X, data is represented by *tokens*. The concept of token can be specialized for each model of computation. For instance, in a discrete event model, tokens may have a value and a time-stamp, while in a data-flow model, they carry a value only. The type of the value which is carried by a token is not taken into account by the MoC, which is only in charge of delivering the tokens by interpreting the relations between the blocks.

The behavior of a block can be described either using a formalism which is external to our framework (for instance in C or Java), yielding an *atomic block*, or by a ModHel'X model. To handle the latter case, we have introduced a special type of block called an *interface block*, which implements hierarchical heterogeneity: the internal model of an interface block may obey a MoC which is different from the MoC of the external model in which the block is used. Interface blocks are a key notion in our framework since they are in charge of adapting the semantics of their
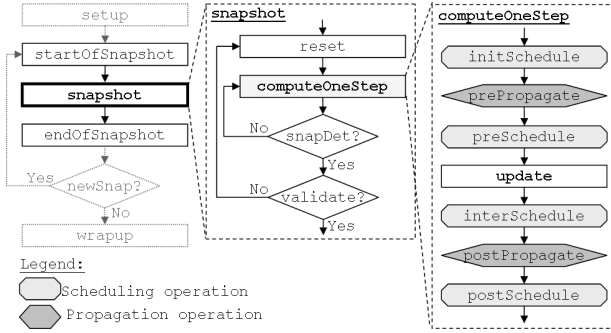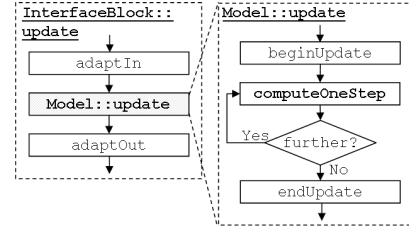
Fig. 2. Generic execution algorithm

Fig. 3. Update on an interface block and its internal model

inner and outer models of computation. They allow the explicit specification of the interactions between different MoCs.

3.4 An imperative semantics for MoCs and their interactions

Computing a snapshot of an heterogeneous model requires to compute the observation of all its parts, which may use different MoCs i.e. different notions of time, control or data. The issue of the consistency of such an observation is similar to the definition of the state of a distributed system [6]. In ModHel'X, we have chosen to define a model of computation as an algorithm for computing observations of the model to which it is associated. For each observation, the algorithm asks the blocks of the model to *update* the state of their interface. The results of the update (output data) are propagated to other blocks by *propagation* operations. We want our execution engine to be deterministic, therefore we observe the blocks sequentially. To ensure the consistency of the computed behavior with the control and concurrency notions of the original model, the MoC must include *scheduling* operations which determine the order in which to update the blocks.

Figure 2 represents the generic structure of our algorithm. This structure is a fixed frame which "standardizes" the way MoCs can be expressed in ModHel'X, but the content of its elements is left free. Therefore, for each MoC, the semantics of the operations of this algorithm has to be described, using an imperative syntax, in order to define the scheduling and propagation "policies" specific to the MoC (non necessary operations can be left empty). The left part of the figure shows the loop which computes the succession of snapshots of the execution of the model. In the computation of a snapshot, the computation of an observation of one block brings into play the scheduling and propagation operations mentioned above and is called a *step* (represented on the right part of Figure 2 under the name `computeOneStep`). The algorithm loops on successive steps until the snapshot is entirely determined (i.e., for most MoCs, when the state of all outputs of the executed model is known). A same block may be updated several times in this loop, what allows the use of non-strict [17] blocks for the computation of fixed point behaviors. Therefore, ModHel'X supports MoCs in which cyclic dependencies are allowed.

The execution of a model traverses the hierarchy thanks to the delegation of the operations of interface blocks to their internal model. Snapshots are realized only at the top level, which represents the whole system. An internal model is only asked to provide a coherent view of its behavior when its interface block is

6

updated. The `update` operations of interface blocks and models are shown on Figure 3. The `adaptIn` and `adaptOut` operations of an interface block allow the modeler to specify explicitly how the semantics of the internal and the external MoCs are adapted before and after the update of its internal model. This may include the adaptation of data, time or control. The `beginUpdate` and `endUpdate` operations in the update of a model are used respectively to take new inputs from the interface block into account, and to provide it with newly determined outputs. The observation of a model may be partial (if it models a non-strict component). The loop which computes the observation must stop when the `further` operation indicates that no more outputs can be determined.

3.5 Implementation and validation

We have experimented our approach in a prototype of ModHel'X based on the Eclipse EMF framework [8]. We use the ImperativeOCL [20] language, an imperative extension of OCL, for describing the semantics of the operations of our algorithm. No interpreter being available for the moment, we translate it into Java. We have successfully implemented several MoCs, such as Finite State Machines (FSM), Discrete Events (DE) and *charts [11]. We are developing a library of MoCs in order to further the validation of our approach. In particular, we are currently working on the UML Statecharts and the Synchronous Dataflow (SDF) MoCs.

3.6 Example of a multi-formalism model in ModHel'X

To illustrate our approach, and in particular the semantic adaptation between a timed and an untimed MoC, we consider a simple hierarchical and heterogeneous model of a coffee machine which works as follows: first the user inserts a coin, then he presses the "coffee" button to get his coffee after some preparation time.

In this model, we take into account the date of the interactions between the user and the machine: insert a coin, push a button, deliver coffee. Therefore, we use the Discrete Events (DE) MoC, which is implemented by SimEvents (The MathWorks), VHDL or Verilog for instance. We represent our user by an atomic block, whose behavior is written in Java. We model the coffee machine as an automaton (with UML Statecharts for instance), because at this stage of the design process, we focus on the logic of its behavior. We consider here a simple version of this MoC called FSM (Finite State Machines), which is similar to the one presented in [19]. Figure 4 shows the global model resulting from the combination of the DE and FSM models. Such a combination is a classical example, which is well addressed by tools like Ptolemy. However, we will see that it is possible to handle the interactions between DE and FSM differently with ModHel'X.

The representation of the structure of the DE model in ModHel'X is straightforward. The representation of the FSM model is more involved because a transition may have two associated behaviors: the evaluation of its guard and its action. Since blocks are the basic units of behavior in ModHel'X, a transition is represented using a block for its guard linked to a block for its action. Relations between guards represent the states.

In DE, when a snapshot is taken, the current time is determined according to
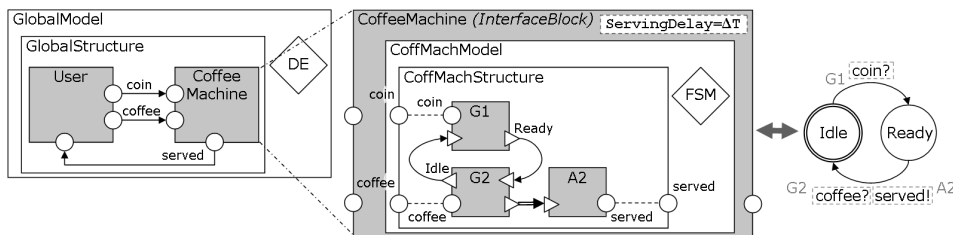
Fig. 4. Global model of the coffee machine and coffee machine automaton

```
DiscreteEventMoC::initSchedule(model:Model){
  List(Block) l := self.timeConstraints
    ->select(t:TimeConstraint | t.time = self.currentTime);
  l.add(self.activeEventList
    ->collect(e:Event | e.destPort.owningBlock));
  // The findARoot function realizes a topological sort
  self.currentComponent := findARoot(l);
}
```

Fig. 5. `initSchedule` operation of the DE MoC

```
CoffeeMachine::adaptOut(){
  if (self.model.structure.getPin("served").storedTokens.
  notEmpty())
    self.getPin("served").storedTokens.add(
      new DEEvent(true, self.coffeeDate
      + self.getParameter("ServingDelay")));
}
```

Fig. 6. Coffee machine `adaptOut` operation

the time stamps of the input events and on the time constraints produced by the blocks. At each computation step, we consider the blocks which have posted a time constraint for the current time and the blocks which are the target of events. We update a minimal element of these blocks according to a topological sort in order to respect causal dependencies. Figure 5 shows the code of the `initSchedule` operation for the DE MoC. This is the only scheduling operation for this MoC, the others (`pre`, `inter` and `postSchedule`) being left empty.

DE and FSM share the notion of event. However, FSM has no notion of time attached to events. So, when a DE event enters FSM, the interface block has to remove its time stamp to make it look like an FSM event. When an FSM event enters DE, the interface block has to give it the "right" time stamp. An acceptable way to proceed is to give it the same time stamp as the most recent incoming event (in particular, this is what is done by Ptolemy). We provide an interaction pattern which realizes this adaptation. However, for our coffee machine, this behavior does not model the serving delay, which is an important characteristic of the model. Therefore, we add a `ServingDelay` parameter to the coffee machine and we modify the pattern so that the time stamp of the `served` event is the time stamp of the `coffee` event plus the `ServingDelay` (see the `adaptOut` operation on Figure 6).

4 Discussion

4.1 Intended workflow and required effort for using ModHel'X

There are two prerequisites to the use of the ModHel'X framework. First, an expert of a modeling language has to describe the structural and semantic elements of this language using our meta-model and our imperative syntax. Since our goal is not to replace existing modeling tools, this expert also defines transformations from the original meta-model of the language to our generic meta-model. This is the difficult part of the work because the semantics of modeling tools is often known intuitively, through the experience we have of the tools. Second, for each pair of MoCs that may interact in heterogeneous models, experts should define interaction patterns, which model standard ways of combining models that obey these MoCs. These

steps represent the main effort needed to benefit from the ModHel'X approach. However, they are done once and for all for each modeling language. Then, system designers can assemble heterogeneous models of the parts of the system, and use these patterns to specify how the models are glued together. Parameters of the patterns allow the designers to fine tune the semantic adaptation at the boundary of two models of computation.

4.2 Supported models of computation

Considering that a same structure of model can be interpreted as an automaton or as a discrete event model depending on the MoC which is associated to it can seem somewhat extreme. However, this choice has proven to be powerful since a tool like Ptolemy supports, on this basis, paradigms as different as finite state machines, ordinary differential equations or process networks. Our approach even supports MoCs for continuous behaviors but, because we address the digital execution of models, we compute a series of discrete observations of such models. This computed behavior is an approximation of the continuous behavior of the system.

4.3 Comparison between ModHel'X and Ptolemy

Ptolemy was our main source of inspiration, but we have extended it on several aspects. One of our main contributions is the explicit specification of the interactions between MoCs (see Section 3.4). Moreover, our approach is based on the observation of blocks and not on the triggering of actors. Thanks to this change of paradigm and to the introduction of time constraints, the execution of a ModHel'X model is not necessarily driven by its root level. Finally, the definition of our abstract syntax as a UML meta-model allows us to rely on model transformation tools from the MDE community to exchange models with other tools in the design chain.

5 Conclusion

We have presented an approach to multi-formalism modeling which provides support for the specification of the semantics of a modeling formalism through the concept of model of computation, and which allows the definition of the interactions between heterogeneous parts of a model through a special modeling construct and using an imperative syntax. This approach relies on the black-box and the snapshot paradigms to compute the observable behavior of a model by combining the behaviors observed at the interface of its components. A generic UML meta-model for representing the structure of hierarchical heterogeneous models has been proposed. On this basis, models of computation are described by giving a specific semantics to the primitive operations of a generic algorithm which computes snapshots of models conforming to the proposed meta-model.

We are currently developing the MoC library of our prototype in order to further the validation of our approach. The rigid structure of the execution algorithm of ModHel'X is a first step toward the definition of MoCs in a fixed frame with formal semantics. However, for the moment, our imperative syntax is still too close to Java to have a formal semantics. Future work will also address the verbosity of the

description of the models of computation by defining higher level constructs over our current imperative syntax.

References

[1] Balarin, F., L. Lavagno, C. Passerone, A. L. S. Vincentelli, M. Sgroi and Y. Watanabe, *Modeling and designing heterogeneous systems*, Advances in Concurrency and System Design (2002).

[2] Basu, A., M. Bozga and J. Sifakis, *Modeling heterogeneous real-time systems in BIP*, in: *4th IEEE International Conference on Software Engineering and Formal Methods (SEFM06)*, 2006, pp. 3–12.

[3] Benveniste, A., B. Caillaud, L. P. Carloni and A. L. Sangiovanni-Vincentelli, *Tag machines*, in: *Proceedings of the 5th ACM International Conference On Embedded Software (EMSOFT 2005)* (2005), pp. 255–263.

[4] Bruneton, E., T. Coupaye and J. Stefani, *The fractal component model specification* (2004).

[5] Burch, J. R., R. Passerone and A. L. Sangiovanni-Vincentelli, *Overcoming heterophobia: Modeling concurrency in heterogeneous systems*, in: *Proceedings of the second International Conference on Application of Concurrency to System Design*, 2001, p. 13.

[6] Chandy, K. M. and L. Lamport, *Distributed snapshots: Determining global states of distributed systems*, ACM Transactions on Computer Systems **3** (1985), pp. 63–75.

[7] de Lara, J. and H. Vangheluwe, *ATOM³: A tool for multi-formalism modelling and meta-modelling*, in: *5th Fundamental Approaches to Software Engineering International Conference (FASE 2002)*, 2002, pp. 595–603.

[8] Eclipse Foundation, *Eclipse Modeling Framework (EMF)*.
URL http://www.eclipse.org/modeling/emf/

[9] Eker, J., J. W. Janneck, E. A. Lee, J. Liu, X. Liu, J. Ludvig, S. Neuendorffer, S. Sachs and Y. Xiong, *Taming heterogeneity – the Ptolemy approach*, Proceedings of the IEEE, Special Issue on Modeling and Design of Embedded Software **91** (2003), pp. 127–144.

[10] Fritzson, P. and V. Engelson, *Modelica — A unified object-oriented language for system modeling and simulation*, in: *European Conference on Object-Oriented Programming (ECOOP98)*, 1998, pp. 67–90.

[11] Hardebolle, C., F. Boulanger, D. Marcadet and G. Vidal-Naquet, *A generic execution framework for models of computation*, in: *Proceedings of the 4th International Workshop on Model-based Methodologies for Pervasive and Embedded Software (MOMPES 2007), at the European Joint Conferences on Theory and Practice of Software (ETAPS 2007)*, IEEE Computer Society, 2007, pp. 45–54.

[12] Henzinger, T. A. and J. Sifakis, *The embedded systems design challenge*, in: *Proceedings of the 14th International Symposium on Formal Methods (FM), Lecture Notes in Computer Science*, Springer, 2006, pp. 1–15.

[13] Karsai, G., A. Agrawal, F. Shi and J. Sprinkle, *On the use of graph transformations for the formal specification of model interpreters*, Journal of Universal Computer Science, Special issue on Formal Specification of CBS **9** (2003), pp. 1296–1321.

[14] Lee, E. A. and A. L. Sangiovanni-Vincentelli, *A framework for comparing models of computation*, IEEE Trans. on CAD of Integrated Circuits and Systems **17** (1998), pp. 1217–1229.

[15] Levendovszky, T., L. Lengyel and H. Charaf, *Software Composition with a Multipurpose Modeling and Model Transformation Framework*, in: *IASTED on SE*, Innsbruck, Austria, 2004, pp. 590–594.

[16] Maraninchi, F. and T. Bouhadiba, *42: Programmable models of computation for a component-based approach to heterogeneous embedded systems*, in: *6th ACM International Conference on Generative Programming and Component Engineering (GPCE'07)*, 2007, pp. 1–3.

[17] Meyer, B., "Introduction to the Theory of Programming Languages," Prentice Hall, Hemel Hempstead (U.K.), 1990.

[18] Mosterman, P. J. and H. Vangheluwe, *Computer automated multi-paradigm modeling: An introduction*, Simulation: Transactions of the Society for Modeling and Simulation International **80** (2004), pp. 433–450, special Issue: Grand Challenges for Modeling and Simulation.

[19] Muller, P.-A., F. Fleurey and J.-M. Jézéquel, *Weaving executability into object-oriented meta-languages*, in: *Proceedings of the 8th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems (MODELS/UML 2005)*, 2005, pp. 264–278.

[20] OMG, *Meta Object Facility (MOF) 2.0 Query/View/ Transformation specification* (2005).
URL http://www.omg.org/cgi-bin/doc?ptc/2005-11-01

[21] OMG, *Unified Modeling Language: Infrastructure – version 2.1.1* (2007).