



HAL
open science

Execution Framework of the GEMOC Studio (Tool Demo)

Erwan Bousse, Thomas Degueule, Didier Vojtisek, Tanja Mayerhofer, Julien Deantoni, Benoit Combemale

► **To cite this version:**

Erwan Bousse, Thomas Degueule, Didier Vojtisek, Tanja Mayerhofer, Julien Deantoni, et al.. Execution Framework of the GEMOC Studio (Tool Demo). Proceedings of the 2016 ACM SIGPLAN International Conference on Software Language Engineering, Oct 2016, Amsterdam, Netherlands. pp.8. hal-01355391v1

HAL Id: hal-01355391

<https://inria.hal.science/hal-01355391v1>

Submitted on 16 Sep 2016 (v1), last revised 17 Oct 2016 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Execution Framework of the GEMOC Studio (Tool Demo)

Erwan Bousse

TU Wien, Austria
bousse@big.tuwien.ac.at

Thomas Degueule

Inria, France
thomas.degueule@inria.fr

Didier Vojtisek

Inria, France
didier.vojtisek@inria.fr

Tanja Mayerhofer

TU Wien, Austria
mayerhofer@big.tuwien.ac.at

Julien Deantoni

Université Côte d'Azur, I3S, France
julien.deantoni@polytech.unice.fr

Benoit Combemale

Univ. Rennes 1 and Inria, France
benoit.combemale@inria.fr

Abstract

The development and evolution of an advanced modeling environment for a Domain-Specific Modeling Language (DSML) is a tedious task, which becomes recurrent with the increasing number of DSMLs involved in the development and management of complex software-intensive systems. Recent efforts in language workbenches result in advanced frameworks that automatically provide syntactic tooling such as advanced editors. However, defining the execution semantics of languages and their tooling remains mostly hand crafted. Similarly to editors that share code completion or syntax highlighting, the development of advanced debuggers, animators, and others execution analysis tools shares common facilities, which should be reused among various DSMLs. In this tool demonstration paper, we present the execution framework offered by the GEMOC studio, an Eclipse-based language and modeling workbench. The framework provides a generic interface to plug in different execution engines associated to their specific metalanguages used to define the discrete-event operational semantics of DSMLs. It also integrates generic runtime services that are shared among the approaches used to implement the execution semantics, such as graphical animation or omniscient debugging.

Categories and Subject Descriptors D.2.2 [Software Engineering]: Design Tools and Techniques; I.6 [Simulation and modeling]: Types of Simulation—Discrete event

Keywords domain-specific modeling language, language and modeling workbenches, model execution, debugging

1. Introduction

The integration of domain-specific concepts and best practices into Domain-Specific Modeling Languages (DSMLs) can significantly improve software and systems engineers productivity and system quality (Hutchinson et al. 2011b). Domain-specific models are used in the development processes to reason and assess specific properties over the system under development as early as possible. This usually leads to a better and cheaper design as more alternatives can be explored. While many models only represent structural aspects of systems, a large amount express behavioral aspects of the same systems. Behavioral models are used in various areas (e.g., enterprise architecture, software and systems engineering, scientific modeling, etc.), with very different underlying formalisms (e.g., business processes, orchestrations, functional chains, scenarios, protocols, activity diagram, etc.).

The development of DSMLs has only been recently recognized as a significant software engineering task in itself (Hutchinson et al. 2011a). To support the tedious task of implementing tool-supported DSMLs, specialized tools (aka. language workbenches) have been proposed during the last decade to provide generic or generated syntactic services for DSMLs. From a DSML specification, Eclipse-based tools, such as Sirius¹ and Xtext (Eysholdt and Behrens 2010), automatically provide advanced services for graphical and textual edition. This includes parsers, syntax highlighting, auto completion or “quick fix” suggestion systems.

For behavioral models, early dynamic validation and verification (V&V) techniques are required to ensure that they are correct with regard to their intended behavior, such as simulation, testing, debugging, or runtime verification. Dynamic V&V techniques may require to execute the models, in accordance with a given specification of the execution semantics of the DSML. Such a specification would eventually provide an execution engine, either in the form of an interpreter (i.e., operational semantics) or a compiler (i.e., translational

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, contact the Owner/Author(s). Request permissions from permissions@acm.org or Publications Dept., ACM, Inc., fax +1 (212) 869-0481.

SLE '16, October 31–November 01, 2016, Amsterdam, Netherlands
Copyright © 2016 held by owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-4447-0/16/10...\$15.00
DOI: <http://dx.doi.org/10.1145/http://dx.doi.org/10.1145/2997364.2997384>

¹<http://www.eclipse.org/sirius/>

semantics). However, the expected *runtime* services on top of the execution engine remain mostly hand crafted and dedicated to the execution engine developed for each new DSML. In particular, graphical animation and interactive debugging are common facilities to observe and control an execution to better understand a behavior or to identify the cause of a defect. Similarly to editors that share code completion or syntax highlighting, the development of advanced debuggers, animators, and others execution analysis tools shares common facilities, which should be reused among various DSMLs.

In this tool demonstration paper, we introduce the *execution framework* of the GEMOC Studio, an Eclipse-based language and modeling workbench. This framework provides a generic interface to integrate different execution engines, themselves associated to their specific metalanguages to define the discrete-event operational semantics of DSMLs. The framework also integrates runtime services that can be shared among the various approaches used to implement the execution semantics. These services include graphical animation, omniscient debugging or trace management. We illustrate the framework with the integration of different engines that are applied on the development of modeling environments for various DSMLs. We demonstrate our approach through the definition of different xDSMLs with different approaches, and the presentation of the resulting environments provided by the modeling workbench to animate and debug models.

2. On Model Execution

Debugging, and dynamic V&V techniques in general, require models to be executable, which is achieved by defining the execution semantics of the executable DSMLs (xDSMLs) used to define them. The execution semantics of an xDSML is mainly implemented either as a compiler (aka. translational semantics) or as an interpreter (aka. operational semantics).

A compiler consists in generating executable code, usually targeting an execution environment that provides the tools of interest for the modeling language under development (*e.g.*, virtual machine, debugger, checkers). While this approach allows language designers to reuse existing, possibly efficient, tools for new modeling languages, it comes with two difficulties. First, the definition of the semantics is given in terms of the targeted execution environment. This makes difficult the definition and understanding of the resulting execution semantics. Second, to be useful to the user of the xDSML there is a need for a back-annotation mechanism to trace back the execution results in terms of the initial model.

An interpreter is a virtual machine for a modeling language in charge of executing any conforming models. The interpreter defines the data structure representing the execution state of a model, and the execution rules in terms of endogenous, possibly in-place, transformations of such an execution state. The execution state can be either an extension of the syntax of the modeling language (*e.g.*, the execution state of a statechart could be specified by an additional collection

containing the current states) or a separate data structure defined by its own metamodel (*e.g.*, the execution state of a Petri net may be represented as a matrix). While the main benefit of this approach is to define the execution semantics directly in terms of the concepts of the modeling language, the main drawback is the necessity for any new xDSML to implement all the tooling based on the execution semantics, *e.g.*, a debugger. Developing such a complex tool for a xDSML remains a difficult and error-prone task.

Various approaches have been investigated in the last decade to implement operational semantics (*e.g.*, (Mayerhofer et al. 2013; Engels et al. 2000; Tatibouët et al. 2014; Combemale et al. 2013)), with their own pros and cons (formality, underlying paradigm, abstraction level, or specific concern such as concurrency). Moreover, despite the specificity of each xDSML and the differences between their semantics, there is a common set of runtime facilities which can be expected for all languages. For instance, one may expect the following debugging facilities: control of the execution (pause, resume, stop), representation of the current state (*i.e.*, model animation), breakpoint definition, step into/over/out and step forward/backward. To support the various approaches to implement the discrete-event operational semantics of xDSMLs, and to drastically reduce the development cost of runtime tools, we present in the next section the generic execution framework provided in the GEMOC studio. This execution framework is used to integrate different approaches to define the operational semantics of xDSMLs (incl., the meta-language and the associated execution engine), and provides various generic facilities for model simulation, graphical animation and debugging.

3. GEMOC Execution Framework

3.1 Overview of the GEMOC Studio

The GEMOC Studio² is an Eclipse package atop the Eclipse Modeling Framework (EMF), which includes:

- The *GEMOC Language Workbench*: to be used by language designers to build and compose new xDSMLs,
- The *GEMOC Modeling Workbench*: to be used by domain designers to create, execute and coordinate models conforming to xDSMLs.

The different concerns of a DSML, as defined with the tools of the language workbench, are automatically deployed into the modeling workbench. They parametrize a generic execution framework that provides various generic services, such as graphical animation, debugging tools, trace and event managers, timeline visualizations, etc.

3.2 Overview of the Execution Framework

Figure 1 shows an overview of the execution framework offered by the GEMOC studio. At the middle, the xDSML de-

²<http://gemoc.org/studio>

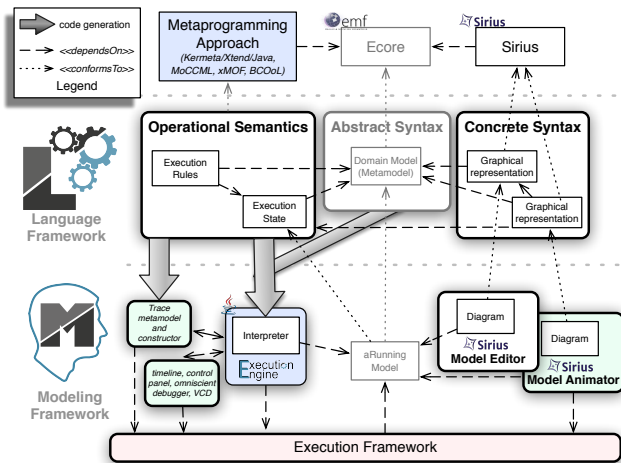


Figure 1: Overview of the GEMOC Execution Framework
 The diagram illustrates the GEMOC Execution Framework, which is composed of several layers and components. At the top, the **Metaprogramming Approach** (supporting languages like Xtext, Xpascal, MoCCML, xMOF, BCool) is linked to **Ecore** and **Sirius**. Below this is the **Language Framework**, which includes **Operational Semantics** (with Execution Rules and Execution State), **Abstract Syntax** (with Domain Model/Metamodel), and **Concrete Syntax** (with Graphical representations). The **Modeling Framework** includes a **Trace metamodel and constructor**, an **Interpreter** (with Execution Engine), **aRunning Model**, **Diagram**, **Sirius Model Editor**, and **Diagram Model Animator**. At the bottom is the **Execution Framework**. A legend indicates symbols for code generation, dependencies, and conformance.

fin in the language workbench is depicted. It is composed of abstract and concrete syntaxes, and of operational semantics. For a given xDSML, the operational semantics are defined using a specific metaprogramming approach (e.g., specific model transformation languages). Since the GEMOC Studio is based on EMF, Ecore is used to define the abstract syntax, and at runtime the executed model is a set of EMF objects. For defining the concrete syntax, the Sirius toolkit is used. For more information and examples on the language workbench, please refer to the official documentation³.

At the bottom, the modeling workbench supported by the *execution framework* is shown. This workbench allows the user to define an executable model conforming to an xDSML, and execute it using an execution engine and a selection of addons. An *addon* is developed using the execution framework and provides a set of runtime services (e.g., animation or debugging). An *execution engine* is developed using the execution framework and is specific to a metaprogramming approach. It is responsible for integrating the interpreter of the considered xDSML – developed using the same metaprogramming approach as the engine – with the addons provided in the studio. This implies sending notifications to addons regarding the execution (e.g., beginning of the run, start or end of a rule, etc.). By reacting to notifications, an addon may query the engine to ask for information, which can be used to provide a view that gets updated during the execution, or control the execution of the model, or even modify the model. Some addons are generated (e.g., trace management), while others are generic and compatible with any engine. At runtime, the executed model contains a dynamic execution state that is modified by the interpreter and by addons.

3.3 Interface of the Execution Framework

There are many approaches to define discrete-event operational semantics for an xDSML. Each approach has different

characteristics, such as how to initialize the execution state of a model, or how to provide a way to control the application of the semantics. In order to manage all these different situations, the GEMOC execution framework provides an API to define and integrate (discrete-event) execution engines. This API defines an engine as a component with two main operations: *initialize* to load an xDSML and a model, and to prepare the transformation; and *execute* to run the transformation. In addition to these two operations, the API comprises operations to add or remove addons, to access the current stack of execution steps, to get or set the engine status (started, paused, stopped, etc.), to access the execution context (e.g., the model, the execution mode) or to start/stop the engine. This API is used both by the generic part of the framework dedicated to starting an execution, and by the addons which may need to access information or control the execution.

In addition to complying with this API, the integration of an engine in the framework requires that notifications are sent to the attached addons during the execution (i.e., during the *execute* operation). For this purpose, the framework defines the common notion of *execution step*, a step being the application of an execution rule of the operational semantics. Note that an engine implementation may consider that only a subset of the rules lead to execution steps (e.g., by annotating the semantics). When an engine is about to apply such rules, it must create a step object containing information about the executed rule (identifiers of the rule, parameters given, etc.). Then this step object must be used to notify addons of both the start and the end of a step. Step objects are also used in the framework for other purposes, such as the storage of a stack of the steps currently being executed, or to be directly stored within an execution trace.

Since a significant part of the logic is common to all execution engines, the framework provides a basic *abstract execution engine* that can be extended into a concrete engine. This abstract engine implements part of the API described above, such as the services to manage the status, to add or remove addons, and to start or stop. In addition, this abstract engine provides internal services to notify all addons of the progress of the execution, although the task of calling these services at the right instants is left to the concrete engines. In particular, a service called *beforeStep* is provided to be called at the beginning of a step, and second called *afterStep* must be called at the end. Depending on the technique used to define the execution semantics, the integration of these operations in the operational semantics can be done in different ways (instrumentation, event listeners, etc.).

Addons are components that can be defined to provide runtime services. Such services need to be connected to the ongoing execution of a model in order to follow the execution and to extract information (e.g., the content of the execution state) or to control the execution (e.g., pause or provide input). To that effect, the GEMOC execution framework provides an API that defines an addon as a component with at least four

³ http://gemoc.github.io/gemoc-studio/publish/guide/html_single/Guide.html

operations that are synchronously called by the engine during the execution of a model: *engineStarted* when the execution starts, *engineStopped* when the execution ends, *aboutToExecuteStep* when the engine is about to start an execution step, *stepExecuted* when an execution step finished. Within the implementation of one of these operations, an addon can access the engine and its status, the executed model, or even the graphical interface of the studio. Therewith, an addon can accomplish a large diversity of tasks, such as changing the executed model, pausing or stopping the execution, displaying information, or sending some input data to the engine.

4. Integration of Different Execution Engines

Using the API for execution engines that we described, different metaprogramming approaches have been integrated to describe discrete-event operational semantics. In particular, we implemented four main execution engines:

Java Engine. The Java engine is dedicated to operational semantics that are entirely defined using any Java-based language, such as Java, Xtend (Efftinge et al. 2012) or Kermeta (Jézéquel et al. 2015). The assembling of the various DSML concerns (incl., abstract and concrete syntaxes, and operational semantics) is made consistent thanks to Melange (Degueule et al. 2015). The *initialize* operation consists in loading the model, and searching for the entry point specified by the xDSML (i.e., the *main* method of the semantics). The *execute* operation simply consists in starting this entry point. However, to interweave notifications to addons in between execution steps, it is required that the Java semantics themselves notify the engine both before and after the performed steps (e.g., by calling *beforeStep* and *afterStep* provided by the abstract engine), otherwise the engine has no way to be informed of the execution. These calls can be either introduced manually at the beginning and at the end of chosen execution rules (i.e., Java methods), or can be introduced externally using aspect-oriented programming (e.g., AspectJ), or can even be realized using a feature provided by the metalanguage (e.g., the *@Step* annotation provided by Kermeta).

Java+MoCCML Engine. The Java+MoCCML engine is dedicated to operational semantics where the (possibly concurrent and timed) control is described in the MoCCML formal language (Deantoni et al. 2015), while the execution rules are written in any Java based language (e.g., Kermeta) (Combemale et al. 2013). In this case, the *initialize* operation creates a so called concurrency model according to the operational semantics, where some relevant events are defined, constrained together and linked to execution rules defined in Java (or any Java-based language). Then the *execute* operation consists in a loop that asks for the next possible steps to the Timesquare solver (Deantoni and Mallet 2012). There are potentially several possible steps since it considers all acceptable interleavings of the events. Once a step is selected, it calls the corresponding Java methods as required. Notifi-

cations to and from addons are received and sent directly by the engine during this execution loop.

xMOF Engine. The xMOF engine supports the execution of operational semantics defined with xMOF (Mayerhofer et al. 2013). With xMOF the data structure for representing the execution state of a model is defined with a metamodel and the rules of the semantics are defined with UML activities conforming to the fUML standard (Object Management Group 2015). For the execution of models, the xMOF engine relies on the virtual machine of fUML. In particular, the *initialize* operation of the xMOF engine loads the model to be executed as well as an optional input model defining input values to the executed model (e.g., the initial token distribution of a Petri net). These models are then handed over to the fUML virtual machine for execution in the *execute* operation of the xMOF engine. The fUML virtual machine provides sophisticated mechanisms for controlling and observing the execution of a model. These mechanisms are utilized by the xMOF engine to call the operations *beforeStep* and *afterStep* at the appropriate instants.

BCOoL Engine. The engine supports the behavioral coordination of heterogeneous models, based on coordination patterns defined using BCOoL (Vara Larsen et al. 2015). The *execute* operation consists in asking a set of coordinated engines what are the next possible steps and to merge and filter these steps to provide the next possible coordinated steps according to the coordination patterns.

5. Runtime Services

Using the API for addons that we described, we implemented a set of generic runtime services that can be shared among the different execution engines of the GEMOC Studio.

Graphical Animator. The graphical animator is an addon that updates different views in order to display the current execution state of the executed model, hence helping to understand models under execution. Since the GEMOC Studio is based on Eclipse, the animator is connected to the Eclipse debug UI to display the stack of currently executed steps in the *Debug* view, and the values of the execution state in the *Variables* view. In addition, if a graphical concrete syntax was defined for the xDSML using Sirius, the animator updates a graphical representation of the execution state of the model during the execution.

Execution Trace Addons. The execution framework provides two complementary ways to manage execution traces: a generic multibranch trace addon, or a generator of domain-specific multidimensional trace addons. Each trace addon reacts to engine notifications to capture steps and states. The *generic multibranch trace addon* captures generic traces that contain different branches, each linked to a choice made during a non-deterministic situation of the execution. The *generator of domain-specific and multidimensional trace addon* can

automatically produce an addon specific to a given xDSML to create and manage multidimensional traces (Bousse et al. 2015b).

Omniscient Debuggers. Two omniscient debugging addons are provided, each relying on a different sort of trace addon. They can provide services expected by any debugger (e.g., step into/over/out, breakpoints) thanks to an integration with the Eclipse Debug UI. In addition, they provide services to explore the execution *backward* in time (back into/over/out), by relying on execution traces constructed by trace addons. They each provide a view called a *timeline*, which is an interactive graphical representation of the trace. The *multibranch debugger* relies on a multibranch trace to explore non-deterministic and concurrent executions. It allows the creation of new branches by going back in time and making different choices. The *multidimensional omniscient debugger* (Bousse et al. 2015a) relies on a multidimensional trace to explore the different dimensions of the execution, using additional stepwise operations to explore the dimensions.

VCD. The VCD addon provides a representation of the execution in the form of a timing diagram represented in the Value Change Dump format (VCD) defined by the IEEE Standard 1364-1995 and extended in the IEEE Standard 1364-2001. It represents the potentially parallel evolution of the calls of the execution steps. This illustrates the possibility to adapt the GEMOC studio addons to specific domains. In this case it is adapted to the Electronic Design Automation (EDA) domain, which is used to this kind of format.

Stimuli Manager. The stimuli manager is an addon provided to send *stimuli* to an ongoing execution. This addon provides a view showing all the possible stimuli that can be sent to the execution, which depends on the technique used to define the operational semantics, and on the content of the semantics. This addon is interesting for non-deterministic executions that depends on external stimuli, e.g., to simulate stimuli from the external environment.

Step Decider. Lastly, the step decider is an addon provided to make choices in non-deterministic situations during an ongoing execution. It provides a view that shows all the steps possible at a given point in time, and gives the possibility to rely either on an existing decider (e.g., the random decider to make a choice at random), or to manually choose the next step from the displayed possibilities. Like the stimuli manager, this addon is only interesting for non-deterministic semantics with different possible steps at some points in the execution.

6. Related Work

In the last decade, various language workbenches have been proposed with generic or generative approaches to automate the development of *syntactic* services. To name just a few, Xtext and Sirius have been developed on top of EMF to provide textual and graphical advanced editors respectively.

These tools have been developed with a strong connection between the metamodeling environment and the associated services. More recently, a *Language Server Protocol*⁴ has been developed to decouple the IDE and the language syntactic services. This paves the way for using the same services from various IDE that implements this protocol.

In this paper, we propose an execution framework that supports generic and generative approaches for *runtime* services (e.g., graphical animator and omniscient debugging), and a generic API to connect different metaprogramming approaches and associated execution engines. Among many other projects, model execution has been previously investigated within Ptolemy (Eker et al. 2003), a framework where different Model of Computations (MoC) can be integrated based on a unique Java API. While the API can be used to implement different ways of specifying an operational semantics, the framework does not provide any way to communicate from and to user defined addons. Other interesting frameworks for model execution include, ModHel'X (Hardebolle and Boulanger 2008) and ATOM3 (Lara and Vangheluwe 2002), which support multi-formalism modeling and the concurrent execution of heterogeneous models. However, both are bound to their specific metaprogramming approach.

7. Conclusion

Languages workbenches facilitate the development of DSMLs, especially for providing syntactic services. However *runtime* services are mostly handcrafted for each different xDSML and each different metaprogramming approach. Based on a common API, we proposed a framework to integrate any kind of metaprogramming approach used to define discrete-event operational semantics into an execution engine. Notably, implementing this API allows to use and reuse of generic or user-defined runtime services as *addons* that send and receive generic messages to and from the execution engines. Using this framework, we implemented four different engines, each for a specific metaprogramming approach, and a set of runtime services, such as generic debuggers that can be used for any xDSML and engine. As our project is open-source and available online, we are very open to any contributors for implementing additional execution engines (e.g., to support operational semantics defined with other metaprogramming approaches) and additional runtime services.

Acknowledgments

This work is partially supported by the ANR INS Project GEMOC (ANR-12-INSE-0011), the COST Action MPM4CPS (IC1404), the Austrian Science Fund (FWF): P 28519-N31, the Christian Doppler Forschungsgesellschaft CDL-Flex and the BMWF (Austria). The authors also thank Dorian Leroy (Inria) and Cédric Brun (Obéo) for their help in the development of the GEMOC execution framework.

⁴<https://github.com/Microsoft/language-server-protocol>

References

- E. Bousse, J. Corley, B. Combemale, J. Gray, and B. Baudry. Supporting Efficient and Advanced Omniscient Debugging for xDSMLs. In *Proceedings of the 2015 ACM SIGPLAN International Conference on Software Language Engineering (SLE'15)*, pages 137–148. ACM, 2015a. doi: 10.1145/2814251.2814262.
- E. Bousse, T. Mayerhofer, B. Combemale, and B. Baudry. A Generative Approach to Define Rich Domain-Specific Trace Metamodels. In *Proceedings of the 11th European Conference on Modelling Foundations and Applications (ECMFA'15)*, volume 9153 of LNCS, pages 45–61. Springer, 2015b. doi: 10.1007/978-3-319-21151-0_4.
- B. Combemale, J. Deantoni, M. Vara Larsen, F. Mallet, O. Barais, B. Baudry, and R. France. Reifying Concurrency for Executable Metamodeling. In R. F. P. Martin Erwig and E. van Wyk, editors, *6th International Conference on Software Language Engineering (SLE 2013)*, Lecture Notes in Computer Science. Springer-Verlag, 2013. URL <http://hal.inria.fr/hal-00850770>.
- J. Deantoni and F. Mallet. TimeSquare: Treat your Models with Logical Time. In S. N. Carlo A. Furia, editor, *TOOLS - 50th International Conference on Objects, Models, Components, Patterns - 2012*, volume 7304, pages 34–41, Prague, Czech Republic, May 2012. Czech Technical University in Prague, in co-operation with ETH Zurich, Springer. doi: 10.1007/978-3-642-30561-0_4. URL <https://hal.inria.fr/hal-00688590>.
- J. Deantoni, P. Issa Diallo, C. Teodorov, J. Champeau, and B. Combemale. Towards a Meta-Language for the Concurrency Concern in DSLs. In *Design, Automation and Test in Europe Conference and Exhibition (DATE)*, Grenoble, France, Mar. 2015. URL <https://hal.inria.fr/hal-01087442>.
- T. Degueule, B. Combemale, A. Blouin, O. Barais, and J.-M. Jézéquel. Melange: A Meta-language for Modular and Reusable Development of DSLs. In *8th International Conference on Software Language Engineering (SLE 2015)*, Pittsburg, United States, Oct. 2015. ACM. URL <https://hal.inria.fr/hal-01197038>.
- S. Efftinge, M. Eysholdt, J. Köhnlein, S. Zarnekow, R. von Massow, W. Hasselbring, and M. Hanus. Xbase: Implementing Domain-specific Languages for Java. In *11th International Conference on Generative Programming and Component Engineering (GPCE)*, pages 112–121. ACM, 2012. ISBN 978-1-4503-1129-8. doi: 10.1145/2371401.2371419. URL <http://doi.acm.org/10.1145/2371401.2371419>.
- J. Eker, J. W. Janneck, E. A. Lee, J. Liu, X. Liu, J. Ludvig, S. Neuendorffer, S. Sachs, and Y. Xiong. Taming heterogeneity – the Ptolemy approach. *Proc. of the IEEE*, 91(1):127–144, 2003.
- G. Engels, J. H. Hausmann, R. Heckel, and S. Sauer. Dynamic Meta-Modeling: A Graphical Approach to the Operational Semantics of Behavioral Diagrams in UML. In *Proceedings of the Third International Conference on the Unified Modeling Language (UML'00)*, volume 1939 of LNCS, pages 323–337. Springer Berlin Heidelberg, 2000. ISBN 978-3-540-41133-8. doi: 10.1007/3-540-40011-7_23.
- M. Eysholdt and H. Behrens. Xtext: Implement Your Language Faster Than the Quick and Dirty Way. *Object Oriented Programming Systems Languages and Applications*, pages 307–309, 2010. doi: 10.1145/1869542.1869625. URL <http://doi.acm.org/10.1145/1869542.1869625>.
- C. Hardebolle and F. Boulanger. Modhel’x: A component-oriented approach to multi-formalism modeling. In *Models in Software Engineering*, pages 247–258. Springer, 2008.
- J. Hutchinson, M. Rouncefield, and J. Whittle. Model-driven engineering practices in industry. In *2011 33rd International Conference on Software Engineering (ICSE)*, pages 633–642, May 2011a. doi: 10.1145/1985793.1985882.
- J. Hutchinson, J. Whittle, M. Rouncefield, and S. Kristoffersen. Empirical Assessment of MDE in Industry. In *Proceedings of the 33rd International Conference on Software Engineering (ICSE'11)*, pages 471–480. ACM, 2011b. doi: 10.1145/1985793.1985858.
- J.-M. Jézéquel, B. Combemale, O. Barais, M. Monperrus, and F. Fouquet. Mashup of metalanguages and its implementation in the kermeta language workbench. *Software & Systems Modeling*, 14(2):905–920, 2015.
- J. Lara and H. Vangheluwe. Atom3: A tool for multi-formalism and meta-modelling. In R.-D. Kutsche and H. Weber, editors, *Fundamental Approaches to Software Engineering*, volume 2306, chapter Lecture Notes in Computer Science, pages 174–188. Springer Berlin / Heidelberg, 2002. ISBN 978-3-540-43353-8.
- T. Mayerhofer, P. Langer, M. Wimmer, and G. Kappel. xMOF: Executable DSMLs based on fUML. In *Proceedings of the 6th International Conference on Software Language Engineering (SLE'13)*, volume 8225 of LNCS, pages 56–75. Springer, 2013. doi: 10.1007/978-3-319-02654-1_4.
- Object Management Group. Semantics of a Foundational Subset for Executable UML Models (fUML), V 1.2.1, January 2015. <http://www.omg.org/spec/FUML/1.2.1>.
- J. Tatibouët, A. Cuccuru, S. Gérard, and F. Terrier. Formalizing Execution Semantics of UML Profiles with fUML Models. In *Proceedings of the 17th International Conference on Model Driven Engineering Languages and Systems (MODELS'14)*, volume 8767 of LNCS, pages 133–148. Springer, 2014. doi: 10.1007/978-3-319-11653-2_9.
- M. E. Vara Larsen, J. Deantoni, B. Combemale, and F. Mallet. A Behavioral Coordination Operator Language (BCoOL). In *International Conference on Model Driven Engineering Languages and Systems (MODELS)*, number 18, page 462. ACM, Sept. 2015. URL <https://hal.inria.fr/hal-01182773>.