THEME SECTION PAPER

Simulation and analysis of MultEcore multilevel models based on rewriting logic

Alejandro Rodríguez[1] [ORCID] · **Francisco Durán[2]** · **Lars Michael Kristensen[1]**

Abstract

Multilevel modelling (MLM) approaches make it possible for designers and modellers to work with an unlimited number of abstraction levels when specifying domain-specific modelling languages (DSMLs). In this paper, we present a functional infrastructure that allows modellers to define the structure and the operational semantics of multilevel modelling hierarchies, enabling simulation and analysis. Using the MultEcore tool, one can design and distribute the models that compose the language family in a multilevel hierarchy, and specify their behaviour by means of multilevel transformation, so-called multilevel coupled model transformations. We give a rewrite logic semantics to MultEcore's MLM, on which we have based our automated transformation from MultEcore to the rewriting logic language Maude. Then, we rely on Maude to execute MultEcore models and to support analysis techniques, like reachability analysis, bounded and unbounded model checking of invariants and LTL formulas on systems with both finite and infinite reachable state spaces using equational abstraction. We illustrate our developed techniques on a DSML family for Petri nets.

Keywords Multilevel modelling · Domain-specific modelling languages · Model transformations · Verification · Rewriting logic · Maude

1 Introduction

Multilevel modelling (MLM) is a notable research area where models and their specifications can be organised into several levels of abstraction [4]. Indeed, the MLM community has shown that MLM is a favourable approach in domains such as process modelling and software architecture [6,8]. Although there exist diverse approaches for MLM (see [2,17,62,65] for some of them), they all share a common idea: lifting the restriction on the number of levels that designers can use to specify modelling languages.

This restriction is present in traditional Model-Driven Software Engineering (MDSE) approaches which are based on the Object Management Group (OMG) [45] 4-layer architecture such as the Unified Modelling Language (UML) [64] and the Eclipse Modelling Framework (EMF) [46,60]. Like in traditional MDSE approaches, MLM uses abstractions and modelling techniques to tackle the continually increasing complexity of software by considering models as primary artefacts in each phase of the software engineering life-cycle [11]. Using MLM, modellers are no longer forced to fit their modelling language specifications within two levels of abstraction: one for (meta)models and one for their instances. This limitation might indeed be too restrictive for certain situations where the language is large and/or complex, and even more when defining behavioural domain-specific modelling languages (DSMLs). DSMLs that are, for instance, variations on general-purpose languages, i.e., to specify different refinements aimed at specific domains, would require further concretisations of their metamodels. Moreover, these limitations may lead to complications like model convolution, accidental complexity, and mixing concepts belonging to different domains (see, e.g., [6,7,19] for discussions on these issues).

✉ Alejandro Rodríguez
arte@hvl.no

Francisco Durán
duran@lcc.uma.es

Lars Michael Kristensen
Lars.Michael.Kristensen@hvl.no

[1] Western Norway Univ. of Applied Sciences, Bergen, Norway

[2] ITIS Software, University of Málaga, Málaga, Spain

One of the most prominent applications of MDSE is the construction of DSMLs [46]. These are modelling languages that are tailored to a concrete application area [33] which bridge the gap between software engineers and domain experts. DSMLs are usually built on top of a more abstract modelling language, which requires well-defined infrastructures to handle the separation of different abstraction levels. Furthermore, MLM techniques are excellent for the creation of DSMLs, especially when focusing on behavioural languages, since behaviour is usually defined at the metamodel level, while it is executed (at least) two levels below at the instance level [3,18,41]. The reason is that behaviour is reflected in the running instances of the models which in turn conform to their metamodel.

The approach for MLM proposed by the tool MultEcore [39,57], formally specified in [38], rests on the premise that one must be able to specify models (distributed along tree-like hierarchies) which are both generic and precise [41]. Even though various approaches have been proposed for the definition and simulation of behavioural models based on reusable model transformations (e.g., [20,50,52]), these rely on traditional two-level modelling hierarchies. Furthermore, modelling the behaviour through multilevel model transformations [3] and performing execution and analysis in MLM have not been widely explored yet.

Having a hierarchical organisation of the models that are in fact separate artefacts which altogether precisely capture the desired system facilitates future extensions and modifications. This applies not only in the existing levels, but also for adding or removing models to the existing multilevel hierarchy. Therefore, it can help to prevent pollution of models where specialisation of concepts would have to be done in the same model (even if they naturally fit in different levels of abstraction). Furthermore, this enhances modularisation and facilitates extensibility [58].

To cope with execution/simulation of models within the MLM context, Multilevel Coupled Model Transformations (MCMTs) were formally introduced in [41] as a multilevel transformation language that bridges the gap for the execution of multilevel modelling hierarchies. MCMTs are meant to achieve reusable multilevel model transformations for the specification of behaviour. In this paper, we improve the expressive capabilities of MCMTs, with respect to their presentation in [41], by extending them with basic support for attributes, the specification of application conditions to block their execution, and the possibility of expressing multiple patterns through the use of nested parametric boxes.

Even though the potential of the MCMTs has been illustrated in several examples, its practical applicability has been limited. Indeed, the proposal in [41] was only theoretical, and no proper implementation was available. We show in this paper how we have turned the MultEcore editing facilities into a complete development environment in which we cannot only edit our MLM models, but also experiment with them through their simulation and execution, and analyse them by giving access to advanced verification and model checking tools. Having appropriate tools for the development of DSMLs is key, but not our only goal. It is not only about developing models, these models must be *correct*. Having support for the simulation and analysis of such models is key for the use of the MLM techniques. In this paper, we provide a formal semantics of MultEcore models and provide a number of tools for the analysis of such models.

We provide such capabilities for simulation and analysis thanks to a formal specification of MultEcore models in rewriting logic [42,43], and specifically by providing a model-to-model transformation into the rewriting logic language Maude [14,15,21]. As we will see in the rest of the paper, the syntactical facilities of Maude have allowed us to use a representation of MLM hierarchies and MCMTs very close to that of MultEcore. Indeed, this minimal representation distance has facilitated the automation of the bidirectional transformation between them. These transformations give MultEcore users access to the Maude execution engine, which is possibly the most efficient engine for rewriting modulo (combinations of) associativity, commutativity and identity [22,26]. In addition, it also gives access to Maude's formal tool environment, which includes, e.g. tools for reachability analysis, model checking, and confluence and termination analysis [15,23].

This work builds on several previous papers on MultEcore (see, e.g. [38–41,58,68]), but, specifically, it is an extended version of [56]. A detailed discussion on the comparison with this work can be found in Sect. 6. In summary, the main contributions of this paper are: (1) the expressive power of MCMTs have been significantly improved; (2) a formal semantics of MultEcore's MLM hierarchies and behavioural definitions is provided using rewriting logic; (3) formal specifications of MultEcore DSMLs which can be used for simulation and formal analysis; and (4) the MultEcore tool has been extended to give users direct access to such facilities.

All concepts and facilities presented in this paper are illustrated on a case study for a multilevel DSML for Petri nets. Specifically, we show how to design, execute, and verify models. While the case study is described in this paper, we refer the reader to [55] for the complete MultEcore and Maude specifications, including additional details on their analysis.

1.1 Outline

We describe in Sect. 2 the features that characterise our MLM approach using a multilevel DSML for a Petri nets multilevel hierarchy. We level-wise explore each model comprising the hierarchy, from both the structural and behavioural points of

views. Section 3 provides an overview of the infrastructure that transforms the multilevel DSML defined in MultEcore into a Maude specification. This section provides details of the generated Maude specification. We demonstrate the use of such an infrastructure with a case study in Sect. 4, where we perform execution and analysis of a Petri net model of a gas station. In Sect. 5, we discuss related work. Finally, Sect. 6 concludes the paper and outlines directions for future work.

2 Multilevel modelling of Petri nets

In this section, the MultEcore tool and its modelling facilities are summarised. Instead of presenting the previous state of the tool as preliminaries, and then introducing its new features, this section presents the current state of the tool, emphasising on the new features. Although the new features are explained in some more detail, the purpose of the section is to give a complete overview of the facilities of the language. To illustrate the different concepts and techniques discussed in this paper, we use as case study a DSML for Petri nets. In the next sections, we describe each of the models that constitute the Petri net multilevel hierarchy. We depict in Appendix A (Fig. 18) the complete developed PNs multilevel hierarchy (where we omit Ecore at the top).

2.1 The MultEcore tool

The MultEcore tool is designed as a set of Eclipse plugins, giving access to its mature ecosystem (integration with EMF) and incorporating the flexibility of MLM. In the MultEcore approach [41], the abstract syntax is provided by MLM models and the behaviour is provided by Multilevel Coupled Model Transformations (MCMTs) [38,41]. Using the MultEcore tool, modellers can (i) define MLM models using the model graphical editor; (ii) define MCMTs using its rule editor; and (iii) execute and analyse specific models. The execution of MultEcore models rely on a transformation of the models into Maude [15] specifications. When we design a multilevel DSML, we first define its syntax/structure with multilevel modelling hierarchies. Then the behaviour is specified via our multilevel transformation language.

For implementation reasons, MultEcore prescribes the use of Ecore [60] as root graph at level 0 in all example hierarchies. Models are distributed in *multilevel modelling hierarchies*. A multilevel modelling hierarchy in our context is a tree-shaped hierarchy of models with a single root typically depicted at the top of the hierarchy tree. Thus, hierarchies enclose a set of models connected via typing relations. Levels are indexed with increasing natural numbers starting from the uppermost one, having index 0.


Fig. 1 Simple Petri net model

2.2 Petri nets metamodel

Petri nets (PNs) is a well-established formalism to model concurrent systems [47,48]. There is a rich body of theoretical results enabling analysis of PNs, and an enormous set of supporting tools.

A PN is a directed bipartite graph, in which the nodes represent transitions (i.e. events that may occur, represented by rectangles/bars) and places (i.e. states, represented by circles). For example, Fig. 1 shows a Petri net model using a well-known concrete syntax. In it, we find places p1…p4 and transitions tr1 and tr2, where p1 and p2 are connected to tr1 via input arcs, p3 is connected to tr1 through an output arc and to tr2 via an input arc, and finally p4 is an output place of tr2.

The nodes and arcs constitute the static structure of a PN. The dynamic behaviour of the net is given by the *token game*, representing various states of the system. This token game is based on the firing of transitions that leads to the consumption/production of tokens; each fired transition produces a new model state.

A particular state is a snapshot of the system's behaviour. The state of a place is called its *marking*, represented by the presence or absence of tokens—commonly represented as black dots—in the places. In the example shown in Fig. 1 there are three tokens in p1 and two tokens in p2. The current state of the modelled system (marking) is given by the number of tokens in each place.

The increasing complexity of systems has promoted a proliferation of Petri nets variants and extensions during the last decades, as often classical Petri nets are too basic to capture the needs of certain domains. A brief comparison of different kinds of Petri nets can be found in [9]. Although our hierarchy could include other types of PNs, here we only include classical or regular PNs and reset/inhibitor nets [66].

We show in Fig. 2 a PN metamodel aimed to capture the abstract concepts of Petri nets. This metamodel represents the level 1 of the hierarchy (Fig. 18a). The purpose of this model is merely structural. In other words, subsequent levels below it should define the concrete semantics of the PN language(s) (as we show in this section). Thus, as we can see in the figure, a PN contains nodes, each of which can be either a Place or a

Fig. 2 Conceptual Petri nets metamodel (also shown in Fig. 18a)
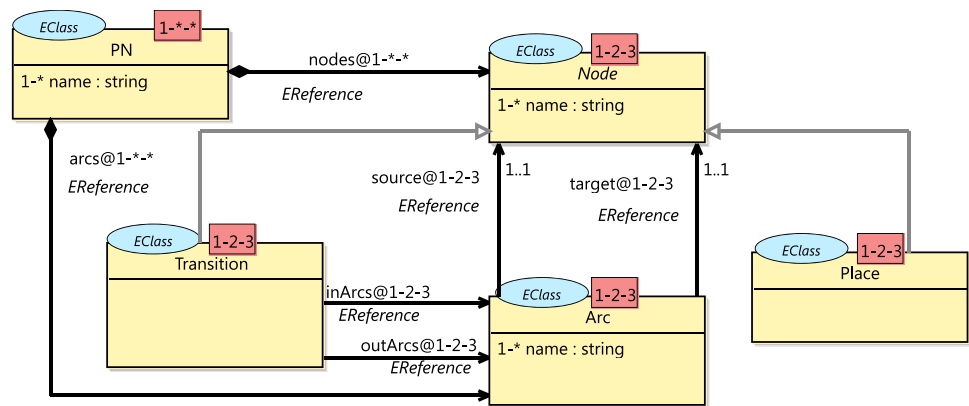
Transition, and Arcs. Note the use of the inheritance relation and the definition of the Node class as abstract class (note the italics), indicating that it cannot be instantiated. As shown in the figure, the type of a node, provided by some element in an upper-level metamodel, is indicated in a (light blue) ellipse at its top left side, e.g. EClass is the type of PN, Node, Transition, Place, and Arc. The type of an arrow is written near the arrow in italic font type, e.g. EReference for arcs, nodes, source, target, inArcs and outArcs. Classes may have attribute declarations, that can be typed by one of the four basic Ecore types, namely Integer, Real, Boolean and String. These attributes can be instantiated in a lower level with a value, as will be illustrated in Sect. 2.5. For the manipulation of attribute values, and the specification of rule conditions, a subset of OCL [12,13,67] is currently supported.[1]

The annotations displayed as three numbers in a (red) box at the top right of each node, and concatenated to the name after "@" for every reference, specify their potencies. *Potency* [34] is a well-known concept in MLM and it is used on elements as a way of restricting the levels at which this element may be used to type other elements. By using potencies on elements, we can define the degree of flexibility/restrictiveness we want to allow on the elements of our multilevel hierarchy. The first two values, *start* and *end*, specify the range of levels below, relative to the current level, where the element can be directly instantiated. The third value, *depth*, is used to control the maximum number of times that the element can be transitively instantiated, or re-instantiated, regardless of the levels where this occurs.

Potency in attributes is displayed as two numbers since attributes can be instantiated only once and it does not make sense to create an instance of such instance. The depth on attributes is therefore always 1 and it is not modifiable. In practical terms, only the first two values (start and end) of the potency are available to the user.

For instance, the potency specified for Arc, Node, Transition and Place is 1-2-3, which means that an element can be directly instantiated one and two levels below (levels 2 or 3 in the hierarchy), and such instances can be re-instantiated up to three additional times. This depth is therefore dependent on the value of the type, and the depth of an element must always be strictly less than the depth of its type.

2.3 Regular Petri nets

The regular Petri nets that we consider in this paper are not restricted to the so-called *Ordinary Petri Nets* where input and output arcs consume or produce, respectively, only a single token [30]. We allow natural numbers (arc weights) on arcs so that more than one token can be added/removed at a time. Following the PNs convention, we denote this number as the *weight* of the arc.

2.3.1 A metamodel for regular Petri nets

Figure 3 shows the *regular-petri-nets* model. It is located at level 2 of the hierarchy (Fig. 18) where we instantiate the concepts defined at level 1. Thus, in this model, we provide the structural basis for the modeller to be able to define further Petri net instance models. InputArc and OutputArc connect regular places and regular transitions. A RegularPlace controls how many tokens it is holding via the numTokens attribute.[2] The weight of arcs is represented as attributes weight of type int in classes InputArc and OutputArc.

2.3.2 Operational semantics of regular Petri nets

The multilevel transformation language that MCMTs define allows us to exploit multilevel capabilities and is powerful enough to specify behavioural descriptions in an operational

---

[1] OCL was chosen since it has been part of UML for several years, is one of the most used languages in EMF-based applications, and it is considered a standard in the MDSE community.

[2] Note that the number of tokens may be calculated with the OCL expression rp.tokens->size(). The attribute is however used to speed up calculations and to illustrate the use of attributes.
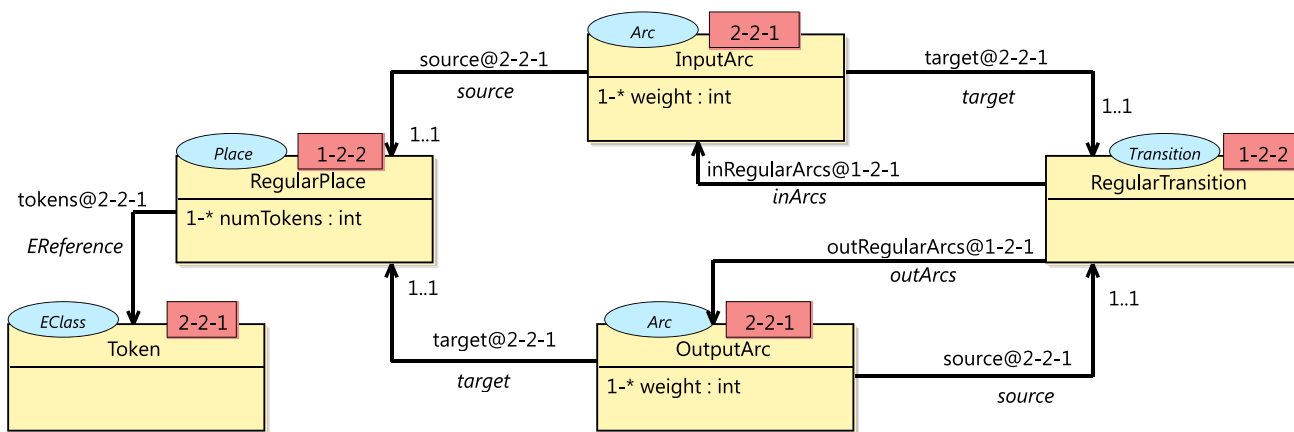
Fig. 3 Regular Petri nets metamodel (also shown in Fig. 18b)

way. Transformation rules can be used to represent actions that may happen in the system. A rule has the form of *LHS* ⇒ *RHS if C*, where *LHS* is a multilevel model pattern (which may contain variables), and *RHS* is a model pattern in which we can use the variables already appearing in *LHS*. *C* is a boolean condition, in which we can use variables from *LHS*. Given a model *M* that represents a state of the system, we say that there is a match of *LHS* on *M* if there is a submodel $M|_p$ of *M* such that for some assignment σ of the variables in *LHS*, we have $LHS\sigma = M|_p$, where $LHS\sigma$ denotes the application of the assignment σ to *LHS*. Given a match of *LHS* on *M*, for some assignment σ and the submodel $M|_p$, the condition $C\sigma$ is evaluated, where, similarly, $C\sigma$ denotes the application of the assignment σ to the condition *C*. If it evaluates to *true*, then the application of the rule consists of the replacement of the submodel $M|_p$ of *M* by $RHS\sigma$, which we denote $M[RHS\sigma]_p$. In other words, if there is a match of the rule on the model, and its condition is satisfied, then the matched submodel is replaced by the model specified in the right-hand side of the rule.

The way to express the behaviour of systems using transformation rules is by specifying rules modelling each of the possible actions that may occur. In PNs, actions occur when transitions are fired. In a regular PN, we only have regular arcs connecting places with transitions. Although transitions can have an arbitrary number of input and output places, such an action can be specified with the MCMT rule *Fire regular transition* depicted in Fig. 4. Specifically, this rule models a transition being fired, taking into account the information of the input places (arcs connected to it) and the output places (arcs where new tokens are to be produced). The FROM and TO blocks describe the left pattern and the right pattern of the rule, respectively. The META block shows a multilevel pattern allowing us to locate types at any level that can be used as individual types for the items in the FROM and TO blocks, respectively. Notice that the META block facilitates

the definition of an entire multilevel pattern, therefore, we can specify several levels within the block.

At the top level of Fig. 4, we mirror parts of the *petri-nets-concepts* model (depicted in Fig. 2), defining elements like Node, Arc, Transition, Place, source, target, inArcs and outArcs as constants—constant elements have their names underlined and their types are not specified. A constant node in an MCMT rule is a *literal value*. It can only match to a node in the hierarchy with the same name in the matched level. In a constant edge, it is the name, the name of the source and the name of the target which must be the same in order to have a match. The use of constants constrains the matching process, significantly reducing the amount of matches. Note that variables like InArc in the lower level of the META block or tk1 in the FROM section, may be instantiated to any element in the MLM model. We depict in the META block those elements and their relationships that are useful for the specification of the FROM and TO patterns.

In the second level of the META block (below the red horizontal line), we capture elements to serve as types to scope the execution of regular PNs. In this level, we find elements as variables. For example, Token, whose type is denoted in the ellipse right above it (EClass). We also have RegPlace, of type Place, and the reference rptk, of type EReference. Similarly, we express attributes (such as numTokens) that will later be used in the levels below. When the element is a variable, the match is based only on finding the right structure in the model. This opens for the definition of generic rules which are applicable to structurally similar hierarchies in which the elements have different names.

Note that the horizontal lines do not enforce consecutiveness between the levels specified in the rule with respect to the hierarchy. This leads to a more natural way of defining that a type is defined at some level above, without explicitly stating at which level. In fact, this also promotes flexibility in case of future modifications of the number of branches
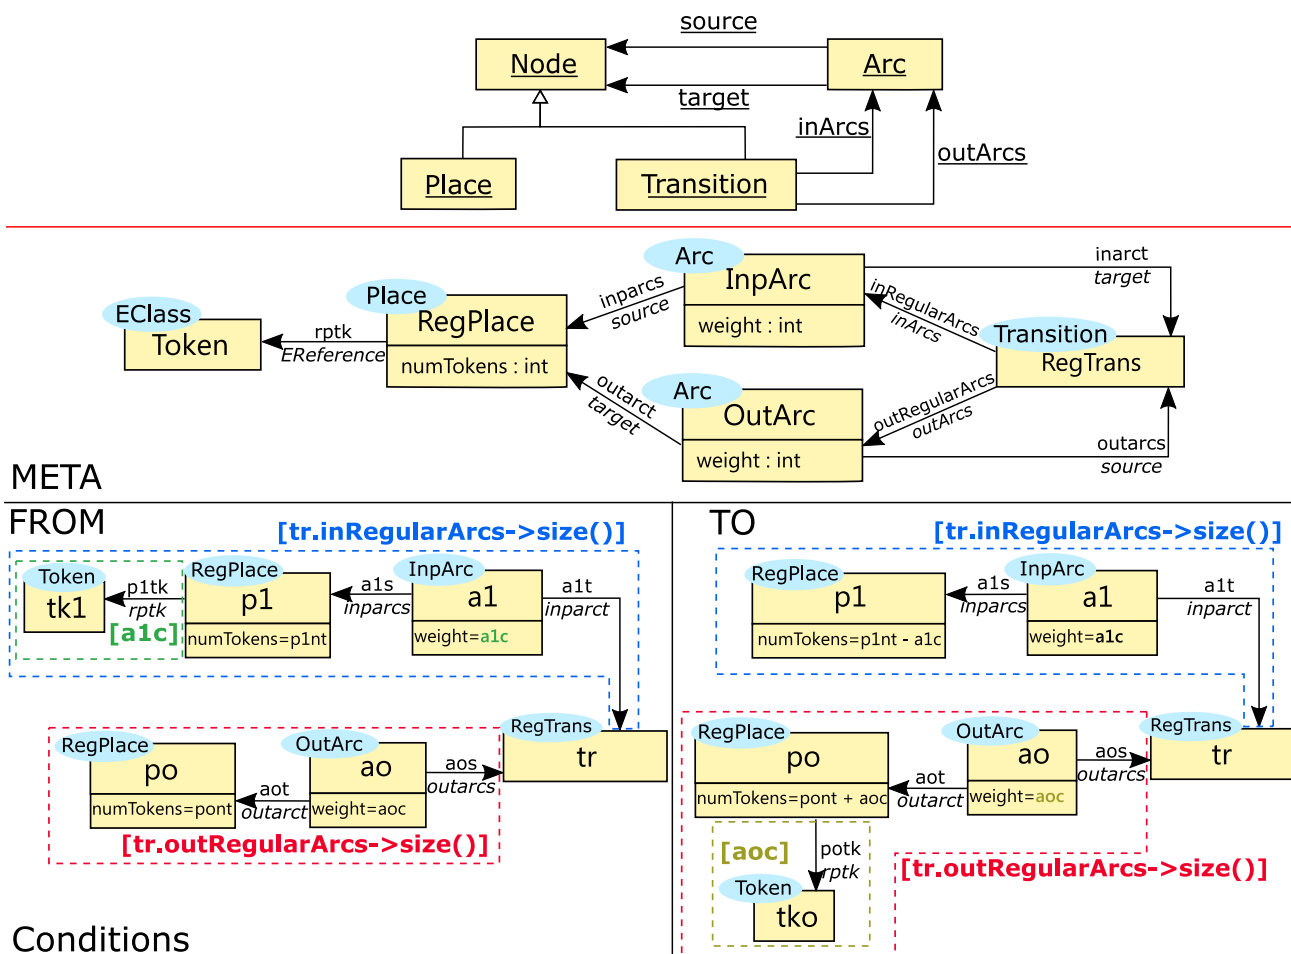
Fig. 4 Rule Fire regular transition: It removes tokens in the input places and creates new ones in the output places

(horizontal dimension) and the depth (vertical dimension) of hierarchies. For instance, the three levels depicted in the rule in Fig. 4 would match to levels 1, 2 and 4 in the multilevel hierarchy depicted in Fig. 18. As the aim of the PNs multilevel hierarchy used as an example here is not to highlight the horizontal/vertical flexibility, we refer the interested reader to [56, Section 4.2] for details on this.

We specify in the FROM block what elements must be found in the model in order to be able to fire a transition. As one can observe, dashed boxes are specified around certain parts of the FROM model. A key point when defining model transformation rules is to make them as reusable as possible. Furthermore, in a PN, there might not only be as many input/output places connected to a transition as one requires, but also an arbitrary number of tokens residing within each of these places. Clearly, it is not practical to define one rule per possible combination of these connections, as the number of rules would rapidly blow up. MCMTs allow the use of nesting boxes to define patterns where its unfolding would

result in a collection of elements. As seen in the rule in Fig. 4, boxes may appear in both sides, and they can be nested.

The blue dashed box in the FROM section of the rule in Fig. 4 encapsulates the nodes tk1, p1 and a1, as well as the references p1tk, a1s and a1t, covering all the potential input places connected to the transition (matched to tr) in the model. The number of instances of this pattern submodel is given by the OCL expression tr.inRegularArcs→size(), also in blue, which represents the number of incoming arcs, i.e. the size of the collection of incoming regular arcs of the transition tr.

In OCL, the size() operator calculates the size of the collection it is applied on. The tr.inRegularArcs expression returns the collection of edges whose source is tr and its type is inRegularArcs. Note, however, that the way in which types are used in MLM is a bit different than for standard OCL. This allows transitive typing, which as we will see below, may be very useful. If instead, as in the condition of the rule, we use tr.inArcs, then we get the collection of edges of type inArcs or

any of its instances. Note that the expression tr.inArcs→size() = tr.inRegularArcs→size() checks whether all the incoming arcs of a given transition tr are of type inRegularArcs. This means that the rule is only applicable on transitions whose arcs are all *regular*. The number of total input (resp. output) arcs, inArcs (resp. outArcs), must be equal to the number of input (resp. output) regular arcs, inRegularArcs (resp. out-RegularArcs).

Analogously, and using the OCL expression tr.outRegularArcs→size(), a second (red) dashed box allows us to specify a number of output places (and corresponding arcs) connected to the transition.

Note the (green) nested box in the FROM part, inside the (blue) box. This inner box allows us to take an arbitrary number of tokens from the input place. For a specific instantiation of the rule, the cardinality of the box is matched to the variable a1c that takes the value of the weight attribute of arc a1. Indeed, given these boxes, a transition may have multiple incoming arcs, and for each incoming place-arc, multiple tokens.

There are also boxes on the TO part. Notice that the input and output arcs are left unmodified, but the appropriate number of tokens are added to the corresponding output places. The number of tokens to put in an output place is provided by the weight attribute of the outcoming arc. The nested (green) box in the TO part, inside the (red) box, indicates that the number of tokens (tko) to be added to each output place po connected to tr via ao, is given by the value aoc of the weight of the arc ao. Finally, note the use of OCL expressions for the manipulation of attributes. In this case, the numTokens attributes of places p1 and po are correspondingly updated: each input place p1 from which some tokens are removed and each output place po that receives tokens, gets its numTokens attribute, respectively, decreased (numTokens = p1nt - a1c) or increased (numTokens = pont + aoc) with the corresponding number of tokens.

In summary, the rule can be executed if the unfolded number of elements is found during the matching process, and all the conditions are satisfied. If this is fulfilled, the model in the TO part is produced. In this case, the execution of the rule removes all the tokens present in each of the input places as specified in the boxes, and creates new tokens on the output places.

2.4 Reset/inhibitor Petri nets

A reset/inhibitor PN [66] is a PN that in addition to regular arcs may also have reset and inhibitor arcs. A reset arc is an input arc that connects a place to a transition and which removes all the tokens of the place when the transition is fired. This is useful as a "cleaning mechanism" in models that capture, e.g. certain environments where messages might be retransmitted and buffers could accumulate old messages.


Fig. 5 Concrete syntax of a reset/inhibitor Petri net example

An inhibitor arc is an input arc which is used to reverse the logic of an input place. With an inhibitor arc, the absence of a token in the input place is what enables the connected transition (not its presence). For instance, inhibitor arcs can be used to delay certain actions until a system is idle, or to wait until the end of a loop.

Figure 5 shows a very simple example of a reset/inhibitor PN in which we have one arc of each type. In this example, p1 is connected to tr1 via a regular input arc (defined in Fig. 3), p2 via a reset arc (denoted with double arrow heads) and p3 via an inhibitor arc (distinguished with a small circle instead of an arrow head). Thus, this transition could be fired according to the semantics of each of the arcs: since (i) p1 has 3 tokens and its regular input arc requires 2; (ii) p2 does not block the firing of the transition, but it will be emptied by its connected reset arc; and (iii) p3 has 0 tokens which fulfils the enabling semantics of the inhibitor arc.
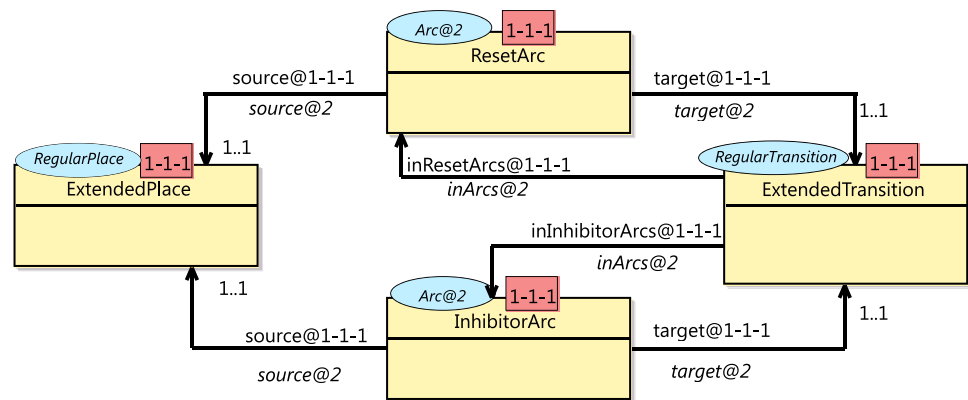
2.4.1 A metamodel for reset/inhibitor Petri nets

Figure 6 shows the model *reset-inhibitor-petri-nets*, placed at level 3 of the hierarchy (see Fig. 18c). The model captures rules extended with the *reset arcs* and *inhibitor arcs*.

As in the model at level 2, the refined Transition in Fig. 6 keeps track of the inhibitor and reset arcs connected to it (through references inInhibitorArcs and inResetArcs, respectively). While ExtendedPlace and ExtendedTransition are typed by elements in the level right above, ResetArc and InhibitorArc nodes are typed directly by Arc, which is located two levels above (as denoted by the @2 after the type).

In our approach, we follow the so-called *abstraction semantics* to organise elements within the multilevel hierarchy based on how abstract they are. Thus, for us, organising elements in different models is a feature that primarily enhances modularisation and promotes separation of concerns [6]. In other words, we do not encourage the *level segregation* principle [35], which establishes that level organisational semantics should be unique, i.e. aligned to one particular organisational scheme, such as *classification* or *generalisation*. Nonetheless, we do encourage the *level cohesion* principle [35], that is, we recommend to organise elements that are semantically close (by means of potency and level organisation).

Fig. 6 Reset/inhibitor Petri nets metamodel (also shown in Fig. 18c)

2.4.2 Behaviour for reset/inhibitor Petri nets

Reset/inhibitor Petri nets have additional semantics that have to be properly managed. The MCMT rule *Fire reset/inhibitor transition* is depicted in Fig. 7. This rule should be compared with the *Fire regular transition* rule shown in Fig. 4. The rule *Fire reset/inhibitor transition* handles the case in which a transition has any number of arcs of any of the three types (regular, reset or inhibitor). In particular, if there are only regular arcs, it behaves as the *Fire regular transition* rule. Observe that the rule in Fig. 7 includes a third META level, where we capture variable elements such as ExtPlace (of type RegPlace), InhArc (representing inhibitor arcs), ResArc (denoting reset arcs) and ExtTrans. As in the levels above, we determine inResetArcs and inInhibitorArcs references with ExtTrans as source, which can be later used in the OCL expressions for the boxes/conditions.

In the FROM block, we need to specify that we might find any number instances of each of the three kinds of arcs. We do it by encapsulating patterns for each of the arc types into a separate box. Corresponding boxes in the right-hand side specify the corresponding action to take on such an arc and its corresponding place. Notice that boxes for regular and reset arcs have corresponding nested boxes specifying the appropriate number of instances. These boxes are described as follows:

> *Regular arcs* The boxes handling regular arcs are exactly as those depicted in the *Fire regular transition* rule, where the box in the FROM block with cardinality tr.inRegularArcs→size() captures each regular arc a1 connecting a place p1 to the transition tr, and removes the number of tokens of each place as given by the weight a1c on the arc. The corresponding number of tokens is then put in the corresponding output places in the TO block.
> *Reset arcs* The box in the FROM block with cardinality tr.inResetArcs→size() captures the reset arcs a2 that connect input places p2 to the transition tr. To remove

all the tokens tk2 present in the connected place p2, the number of tokens in the place is used as cardinality of the inner box. Note that these tokens do not appear in the corresponding box in the TO block. In this way, all of them are removed.

> *Inhibitor arcs* A third box with cardinality tr.inInhibitorArcs→size() captures inhibitor arcs. Since for the transition to be enabled the number of tokens of each place connected via an inhibitor arc must be 0, we simply specify this directly in p3, where it is stated that the attribute numTokens has value zero.

The rest of the rule is similar to what we have already seen.

The condition tr.inArcs → size() = tr.inRegularArcs → size() + tr.inResetArcs → size() + tr.inInhibitorArcs → size() checks that the total number of input arcs is the sum of the number of regular input arcs, the reset arcs and the inhibitor arcs. The condition tr.outArcs→size() = tr.outRegularArcs→size() checks that the total number of output arcs is the number of regular output arcs. These conditions would be key for further extensions of the current PN hierarchy. If the FROM block of the rule matches a submodel of the PN and the conditions are satisfied, the application of the rule results in the removal of the corresponding tokens from the places connected either via regular or reset arcs, and the creation of new tokens in the output places. Notice that the attributes on the places that keep track of the number of tokens get updated.

2.5 Examples of PN models

With the hierarchy described in Sects. 2.2–2.4, we can now define models of regular PNs and models of reset/inhibitor PNs. This is possible, as potency specifications allow us to design the hierarchy in a way where *deep instantiation* [5] can be achieved, being able to instantiate elements residing in any level above.

To illustrate how PNs are represented using the given hierarchy, we show a first example using a concrete syntax for
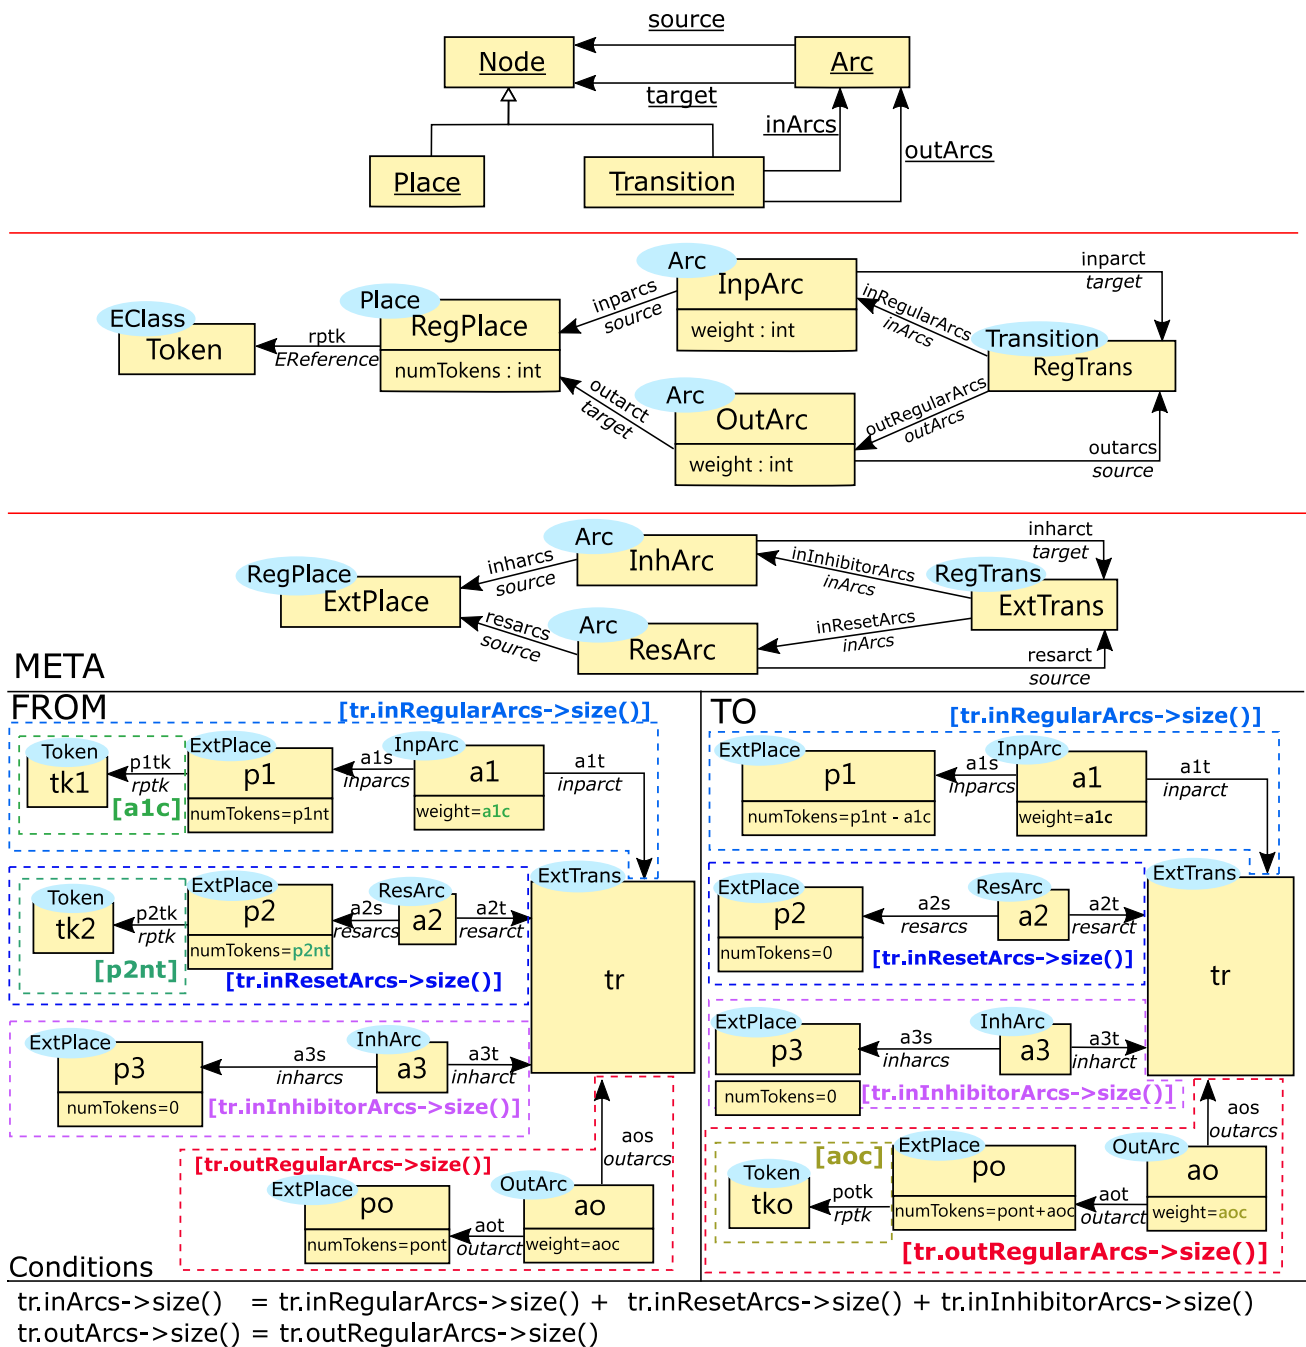
Fig. 7 Rule Fire reset/inhibitor transition: modelling the firing of transitions with regular, reset and inhibitor arcs

Petri nets and then its corresponding one using the MultEcore (abstract) syntax. Figure 8 shows a simple example using regular PNs where four places (two input and two output) and one transition are depicted. To the left we can see that p1 and p2 carry three and two tokens, respectively. Firing tr1 transition would remove 2 tokens from p1 and 1 from p2, and would create 3 and 2 tokens in p3 and p4, respectively, as expressed by the weight in the arcs.

The MultEcore representation of the PN in Fig. 8 is shown in Fig. 9. Since we consider this model at the instance level, we use potency 0-0-0 in the elements. This is used to enforce that elements at the bottom level (in this case level 4) are used purely as instances, which cannot be refined further at levels below it.

As a second example, the MultEcore representation of the PN depicted in Fig. 5 is depicted in Fig. 18d.

Fig. 8 Concrete syntax of a regular Petri net example

3 Formal semantics of multilevel DSMLs

Rewriting logic is used for giving a formal semantics to MultEcore, and specifically, a transformation between MultEcore DSML definitions and Maude specifications provides such a semantics to MultEcore. The Maude representation is therefore key to understand such semantics, and to understand the way MultEcore's DSMLs are executed and analysed. The goal of this section is to introduce Maude and the relationship between MultEcore and Maude, and then the rewriting logic semantics of MultEcore is explained. Specifically, Sect. 3.1 introduces Maude and the use of Maude by the MultEcore tool; Sect. 3.2 presents our Maude semantics of Multilevel hierarchies; and, given its complexity, the formalization of MCMTs is split in Sects. 3.3 and 3.4.

3.1 Maude as a backend tool for MultEcore

Maude [14,15,21] is a specification language based on rewriting logic [42], a logic of change that can naturally deal with states and non-deterministic concurrent computations. A rewrite logic theory is a tuple $(\Sigma; E; R)$, where $(\Sigma; E)$ is an equational theory that specifies the system states as elements of the initial algebra $T_{(6;E)}$, and R is a set of rewrite rules that describe the one-step possible concurrent transi-

tions in the system. Σ is a signature that specifies the type structure (e.g., sorts and subsorts) and operations, and E is the collection of equations and memberships declared in the functional module. Rewrite specifications thus described are executable, if they satisfy restrictions such as termination and confluence of the equational subspecification, and coherence of equations and rules. Maude provides support for rewriting modulo associativity, commutativity and identity, which perfectly captures the evolution of models made up of objects linked by references as in graph grammars.

Although having a formal semantics may be relevant on its own, as we will see in this and the coming sections, a formal semantics of MultEcore's MLM hierarchies and MCMTs is what enables execution of MultEcore models using Maude, and also to carry out different types of analysis on them. The transformation of MultEcore models into Maude specifications, and the execution and analysis results back into the MultEcore tool, is precisely what allows us to provide access to such facilities to MultEcore users. The uses of Maude by the MultEcore tool can be summarised as follows:

Formal specification The Maude specification of multilevel hierarchies and MCMTs represents a formal semantics of MultEcore models in rewriting logic. Based on such formalisation, the transformation MultEcore ⟷ Maude has been automated.

Execution of the specification The Maude specification obtained from MultEcore models using the above transformation is executable, and therefore, it can be used to simulate MultEcore models in Maude. The versatile rewriting engine is not only efficient, but also provides functionalities to customise the way we go through the execution steps. We can simulate our systems by letting
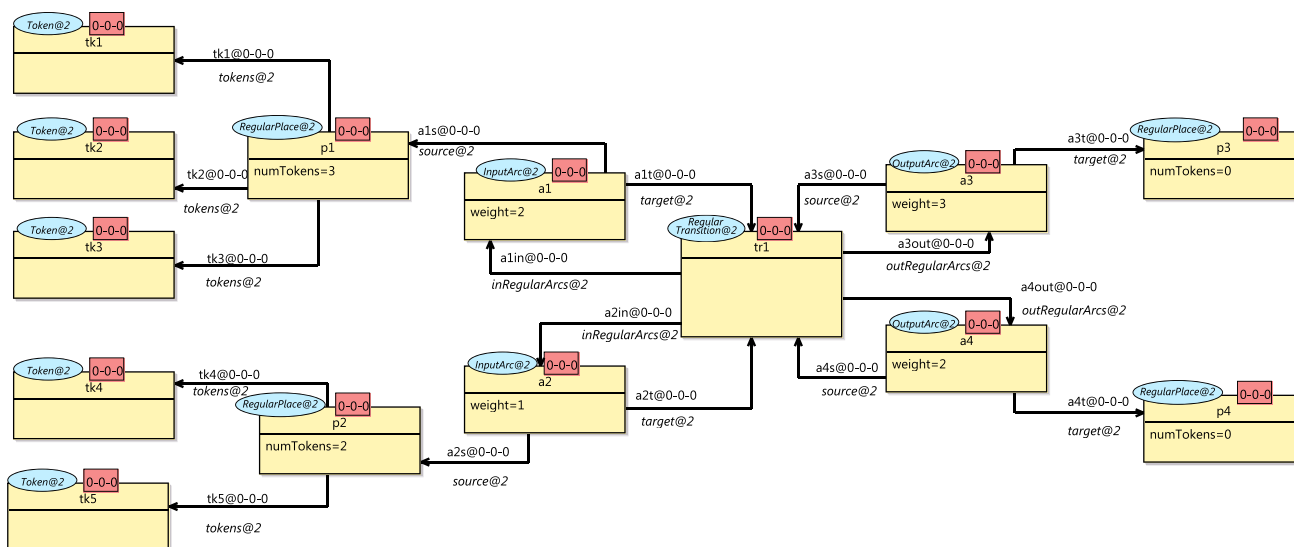

Fig. 9 MultEcore syntax of a regular Petri net example

Maude choose the path to follow, or we can specify a concrete path specifying it step by step, or by means of execution strategies.

Formal environment Once the rewriting logic specification of a MultEcore model is available, we can use the tools in Maude's formal environment to analyse it. For example, we can check properties such as confluence and termination of our specifications, and can also perform reachability analysis, model checking and theorem proving.

The overall MultEcore-Maude infrastructure is outlined in Fig. 10. The left-hand side shows the MultEcore part, where we specify multilevel DSMLs by providing a Multilevel Hierarchy and a set of MCMT rules. The Transformer MultEcore ⟷ Maude has been developed as a bidirectional transformation that takes MultEcore textual specifications and automatically generates Maude specifications, and then takes the XML output files that Maude produces as result of performing execution and analysis, and automatically translates them into MultEcore models graphically displayed.

To grasp an intuition of how the transformation works, each MultEcore object (including both a hierarchy and its MCMTs) is mapped into a corresponding Maude object. References and conditions are handled in exactly the same way, by using references as names, and using the same set of expressions (types and operators) for conditions. The rewriting modulo associativity, commutativity and identity available in Maude captures quite naturally the intended operational semantics of MCMTs. The major challenges were the handling of boxing and the rewriting on multilevel hierarchies. The support for OCL is based on the Maude semantics of OCL proposed in [59].

The right-hand side of Fig. 10 shows the Maude process perspective. The transformer produces a functional module with the equational theory used to represent MLM hierarchies, the MLM Hierarchy, and a system module with rewrite theory that represents the MCMT Rules. The representation of MLM hierarchies and MCMTs is presented in Sects. 3.2–3.4, respectively. We illustrate later (in Sect. 4.1) some of the possibilities for execution and analysis of the models on a case study.

As we will see in this section, MultEcore encapsulates the interaction with the Maude tools, which are hidden to the user. The Maude specification is, however, available to the user, who can interact directly with the Maude environment to get full access to all its features. The complete MultEcore description (both the hierarchy and the MCMTs), the corresponding full Maude specification and the experiments and properties verified can be found in [55].

3.2 A Maude semantics of multilevel hierarchies

In Maude, object-oriented systems can be specified by object-oriented modules in which classes and subclasses are declared, with the usual support for inheritance, and dynamic binding. A class is declared with syntax class $C \mid a_1: S_1, \ldots, a_n: S_n$, where C is the name of the class, a_i are attribute identifiers, and S_i are the sorts of the corresponding attributes. The objects of a class C are record-like structures of the form $< O : C \mid a_1: v_1, \ldots, a_n: v_n >$, where O is the identifier of the object, and v_i are the current values of its attributes.

To represent multilevel metamodels we have introduced declarations to represent multilevel hierarchies as collections of objects each of which represents one of the level models. Specifically, in our approach, a multilevel hierarchy is represented as a structure of sort System of the form
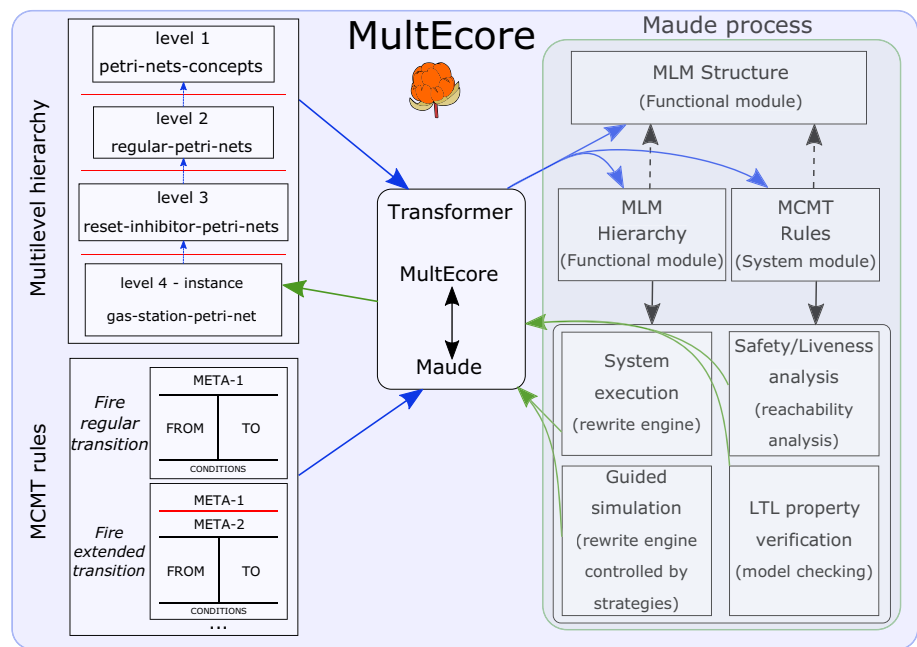
$$\{ \text{model}_1 \ \text{model}_2 \ \ldots \ \text{model}_n \}$$

where each *model_i* is an object of class Model that represents a model in the hierarchy.

Figure 11 shows an excerpt of the Maude specification obtained from the MultEcore Petri net multilevel hierarchy. We have used ellipses (…) to omit some parts due to space limitations. Since levels are numbered starting from 0 (Ecore), the object representing level i's model uses level(i) as identifier. Such an object uses attributes to share the name of the model (name), the name of its immediate metamodel (om), its collection of nodes (elts), and a collection of the relations between these nodes (rels). Elements and relations are themselves represented as objects, of classes Node and Relation, respectively. Each node has attributes to store its name (name), type (type) and its own attributes (attributes). These attributes are again represented as objects with attributes to keep, depending on the level, its name or value (nameOrValue) and its type (type). A relation object has attributes to store its source (source), target (target), and multiplicities, provided by the two usual values (min-mult and max-mult). To avoid name clashes between levels, object identifiers are represented using the operator oid, and nodes and relations using the operator id. Both operators take the level number in which they are defined as first argument, and either a unique number or a string with its actual name.

For instance, the object in lines 1–6 represents level 0, the Ecore model, which has one node with name and type id(0, "EClass") (line 4) and one relation EReference (lines 5–6). Note that the source and target of this relation refer to the names of the source and target nodes, respectively, which in this case is the same id(0, "EClass"). The *petri-nets-concepts* model in lines 7–14 represents the model in Fig. 2 (also Fig. 18a). Lines 10–12 show the representation of node Node, of type EClass, which has several attributes, among which we can see its attribute with name name of type String.

Fig. 10 Infrastructure for the execution and analysis of multilevel modelling hierarchies


```
 1  { < level(0) : Model |
 2        name : "Ecore",
 3        om   : "Ecore",
 4        elts : (< oid(0,  1) : Node     | name : id(0, "EClass"),     type : id(0, "EClass"), attributes : none >),
 5        rels : (< oid(0,  2) : Relation | name : id(0, "EReference"), type : id(0, "EReference")),
 6                                          source : id(0, "EClass"),   target : id(0, "EClass"), ... > >
 7    < level(1) : Model |
 8        name : "petri-nets-concepts",
 9        om   : "Ecore",
10        elts : (< oid(1,  1) : Node     | name : id(1, "Node"), type : id(0, "EClass"),
11                                          attributes : (< oid(1, 2) : Attri | nameOrValue : id(1, "name"),
12                                                                              type : id(1, "String") > )> )>
13                                     ...),
14        rels : (...) >
15    ...
16    < level(4) : Model |
17        name : "reset-inhibitor-petri-net-example",
18        om   : "reset-inhibitor-petri-nets",
19        elts : (...
20             < oid(4, 12) : Node     | name : id(4, "p1"),  type : id(3, "ExtendedPlace"),
21                                       attributes : (< oid(4, 13) : Attri | nameOrValue : 3, type : id(2, "numTokens") >) >
22             ...
23             < oid(4, 28) : Node     | name : id(4, "tk1"), type : id(2, "Token"), attributes : none >
24             ...),
25        rels : (...
26             < oid(4, 36) : Relation | name : id(4, "tk1"),  type : id(2, "tokens"),
27                                       source : id(4, "p1"), target : id(4, "tk1"), ...>
28             ...) > }
```

Fig. 11 Excerpt of the Petri net multilevel hierarchy in Maude representation

The instance model at level 4 is shown in lines 16–28. Note that among its nodes, there is one with name id(4, "p1") (in lines 20-21), of type id(3, "ExtendedPlace")—a node in its metamodel—which has an attribute of type id(2, "numTokens") with value 3.

3.3 Box-free MCMTs in Maude

In Maude, object-oriented systems are axiomatised by equational theories describing their states as algebraic data types and collections of conditional rewrite rules specifying their

behaviour. Rewrite rules are written as

$$\text{crl } [l] : T => T' \text{ if } C$$

where l is the rule's label, T and T' are terms, and C is its guard or condition. As MultEcore's MCMTs, Maude rules describe the local, concurrent transitions that are possible in the system, i.e., when a part of the system state fits the pattern T, then it can be replaced by the corresponding instantiation of T'. Also as for MCMTs, the guard C acts as a blocking pre-

```
1 crl [Fire-1-to-1-Regular-Arcs] :
2    { < level(L1) : Model | name : M,
3        elts : (< O01 : Node | name : id(L1, "Arc"), type : id(0, "EClass"), Atts01 >
4                < O02 : Node | name : id(L1, "Node"), type : id(0, "EClass"), Atts02 >
5                < O03 : Node | name : id(L1, "Transition"), type : id(0, "EClass"), Atts03 >
6                < O04 : Node | name : id(L1, "Place"), type : id(0, "EClass"), Atts04 >
7                Elts),
8        rels : (
9                < O05 : Relation | name : id(L1, "inArcs"), type : id(0, "EReference"), source : id(L1, "Transition"), Atts05 >
10               < O06 : Relation | name : id(L1, "outArcs"), type : id(0, "EReference"), source : id(L1, "Transition"), Atts06 >
11               < O07 : Relation | name : id(L1, "source"), type : id(0, "EReference"), source : id(L1, "Arc"), Atts07 >
12               < O08 : Relation | name : id(L1, "target"), type : id(0, "EReference"), source : id(L1, "Arc"), Atts08 >
13               Rels),
14       Atts >
15   < level(L2) : Model | ... >
16   < level(L3) : Model | name : M''',
17       elts : (< O26 : Node | name : tr_1, type : id(L2, "Transition"), Atts26 >
18               < O27 : Node | name : ao_1, type : id(L2, "OutputArc"),
19                 attributes : (< O28 : Attri | nameOrValue : 1 , type : id(L2, "weight"), Atts28 > Attri27), Atts27 >
20               < O29 : Node | name : po_1, type : id(L2, "Place"),
21                 attributes : (< O30 : Attri | nameOrValue : pont , type : id(L2, "numTokens"), Atts30 > Attri29), Atts29 >
22               < O31 : Node | name : tk1_1, type : id(L2, "Token"), Atts31 >
23               < O32 : Node | name : a1_1, type : id(L2, "InputArc"),
24                 attributes : (< O33 : Attri | nameOrValue : 1 , type : id(L2, "weight"), Atts33 > Attri32), Atts32 >
25               < O34 : Node | name : p1_1, type : id(L2, "Place"),
26                 attributes : (< O35 : Attri | nameOrValue : p1nt , type : id(L2, "numTokens"), Atts35 > Attri34), Atts34 >
27               Elts'''),
28       rels : (< O36 : Relation | name : a1t_1, type : id(L2, "target"), source : a1_1, target : tr_1, Atts36 >
29               < O37 : Relation | name : p1tk_1, type : id(L2, "tokens"), source : p1_1, target : tk1_1, Atts37 >
30               < O38 : Relation | name : a1s_1, type : id(L2, "source"), source : a1_1, target : p1_1, Atts38 >
31               < O39 : Relation | name : aos_1, type : id(L2, "source"), source : ao_1, target : tr_1, Atts39 >
32               < O40 : Relation | name : aot_1, type : id(L2, "target"), source : ao_1, target : po_1, Atts40 >
33               Rels'''),
34       Atts''' >
35   < counter : Counter | value : N >
36   Conf }
37 => { < level(L1) : Model | ... > ---- as in the left-hand side
38     < level(L2) : Model | ... > ---- as in the left-hand side
39     < level(L3) : Model | name : M''',
40       elts : (< O26 : Node | name : tr_1, type : id(L2, "Transition"), Atts26 >
41               < O27 : Node | name : ao_1, type : id(L2, "OutputArc"),
42                 attributes : (< O28 : Attri | nameOrValue : 1 , type : id(L2, "weight"), Atts28 > Attri27), Atts27 >
43               < O29 : Node | name : po_1, type : id(L2, "Place"),
44                 attributes : (< O30 : Attri | nameOrValue : pont + 1 , type : id(L2, "numTokens"), Atts30 > Attri29), Atts29 >
45               < O31 : Node | name : a1_1, type : IA_1,
46                 attributes : (< O32 : Attri | nameOrValue : 1 , type : id(L2, "weight"), Atts32 > Attri31), Atts31 >
47               < O33 : Node | name : p1_1, type : id(L2, "Place"),
48                 attributes : (< O34 : Attri | nameOrValue : p1nt - 1 , type : id(L2, "numTokens"), Atts34 > Attri33), Atts33 >
49               < oid(L3, N) : Node | name : id(L3, N + 1), type : id(L2, "Token"), attributes : none >
50               Elts'''),
51       rels : (< O36 : Relation | name : a1t_1, type : id(L2, "target"), source : a1_1, target : tr_1, Atts36 >
52               < O38 : Relation | name : a1s_1, type : id(L2, "source"), source : a1_1, target : p1_1, Atts48 >
53               < O39 : Relation | name : aos_1, type : id(L2, "source"), source : ao_1, target : tr_1, Atts39 >
54               < O40 : Relation | name : aot_1, type : id(L2, "target"), source : ao_1, target : po_1, Atts40 >
55               < oid(L3, N + 2) : Relation | name : id(L3, N + 3), type : id(L2, "tokens"),
56                                              source : po_1, target : id(L3, N + 1), min-mult : 1, max-mult : 1 >
57               Rels'''),
58       Atts''' >
59   < counter : Counter | value : N + 4 >
60   Conf }
61 if L1 < L2 /\ L2 < L3
62 /\ tr . inArcs -> size() = tr . inRegularArcs -> size() /\ tr . outArcs -> size() = tr . outRegularArcs -> size() .
```

Fig. 12 Excerpt of the Maude rewrite rule corresponding to the box-free version of the *Fire regular transition* MCMT rule

condition: a conditional rule can only be fired if its condition is satisfied. Rules may be given without label or condition.

Given the representation of multilevel hierarchies presented in the previous section, the transformation of MCMT rules without boxes is straightforward. Basically, since the META section does not change, its corresponding representation appears in both left- and right-hand sides. The left-hand side of the rule is completed with the representation of the FROM block, and the representation of the TO block is added to its right-hand side. Variables in MCMTs are represented as Maude variables, and conditions are placed in the conditions as such. Since we restrict conditions to basic types, basic operators and certain selected OCL operations, the expressions can be handled directly by Maude. The last issue to consider is new names and identifiers in right-hand sides. As we explained above, identifiers are represented using the id and oid operators, which take the level in which the object is created and a unique number as arguments. Such numbers are generated using a Counter object whose value attribute gets increased every time a new identifier is created.

Let us illustrate this general procedure on a specific example. Consider the rule *Fire regular transition* depicted in Fig. 4, but let us assume first that it has no boxes in it, that is, let us assume that it takes one single token from the unique input place of a transition and moves it to its unique output place. If this was the case, the corresponding generated Maude rule would be the one shown in Fig. 12. Again, we have used ellipses to shorten the listing.

The left- and right-hand sides of the rule are given in lines 16–34 and 39–58, respectively. The corresponding condition is in lines 61–62. The META section is represented in lines 2–15 and 37–38. The three objects representing the META section are replicated in both sides. They provide the appropriate context for the rule, but it is in the FROM and TO sections where the actual change is modelled. Notice the use of variables to identify the levels (L1, L2 and L3). These are used to match any specific levels in the hierarchy on which the rule is applied. They do not need to be consecutive levels, the only restriction is given in the condition, where it is checked that L1 < L2 < L3 (line 61). Notice that the Maude rule represents quite closely the corresponding MultEcore MCMT. For instance, constants in the MCMT rules are mapped into Maude constants and ground terms. Variables are used both to represent free elements in the rules and also any other elements not explicitly specified. The condition of the MCMT rule is written as such in line 62. Finally, notice the use of the Counter object, which in the left-hand side has some value N (line 35) and in the right-hand side has value N + 4 (line 59) since four new identifiers are introduced.

The model changes applied by the *Fire regular transition* rule have been highlighted by framing them to ease its comprehension (see the rectangles in Fig. 12). Specifically, given a transition tr_1 (line 17) with only one place (notice the

weight 1 of the input arc in lines 23–24), and given one of the tokens in it (the token is specified in line 22 and the relation associating it to the place in line 29), the rule removes such a token and creates a new one (lines 49 and 55–56). The number of tokens in the input place is decremented (lines 26 and 48) and the number of tokens in the output place is incremented (lines 21 and 44). Finally, the created token and relation objects (lines 49 and 55–56) have identifiers oid(L3,N), oid(L3,N+1), oid(L3,N+2), and oid(L3,N+3).

3.4 Meta-programming for MCMTs boxes

As illustrated with the rules depicted in Figs. 4 and 7, boxes allow us to express very general situations in a quite intuitive way. However, Maude does not provide any mechanism similar to that of MCMT boxes, and therefore the transformation is not as simple. To handle boxes in a generic and efficient way, we use Maude's meta-programming capabilities to unfold boxes at runtime as needed.

If we revisit the *Fire regular transition* rule in Fig. 4, this time considering its boxes, we know that each time the rule is applied, depending on the specific situation, there will be a number of replicas of each of the boxes. Actually, notice that we may have multiple boxes on both sides, with different cardinalities, and we can have nested boxes, as many times as needed. Note also that these cardinalities are explicitly specified, otherwise, for example, a given transition could be applied taking an arbitrary number of tokens from several of the available input places. These cardinalities could be provided as OCL expressions, which need to be evaluated to get the corresponding value at the time it is required. In this particular case, the transition has tr.inRegularArcs->size() input regular arcs, each of which has a weight a1c which specifies the number of tokens to be removed from it when the transition is fired. Similarly, the transition has tr.outRegularArcs->size() regular output arcs, which provides the number of output places, each of which has a weight that indicates the number of tokens to be created on that place.

The only assumption that we make to handle boxes is that their cardinality must be greater than zero. In case we want to consider the possibility of zero replications of a box, we need to provide the corresponding rule without such a box. This is the case for the rule in Fig. 7. The cases in which we have transitions with no reset, inhibitor or regular arcs must be handled in different Maude rules. Although these cases can be handled by automatically creating the zero-case corresponding rule, we focus here on the general case.

An MCMT rule with boxes produces two Maude rules (plus the corresponding ones for the zero-cardinality cases). Excerpts of the two rules for the MCMT rule *Fire regular transition* in Fig. 4 are shown in Figs. 13 and 14. The first one of the rules has a left-hand side as if there were no boxes in it. That is, the left-hand side of the rule in Fig. 13 is exactly

as the left-hand side of the rule in Fig. 12. In its condition, the boxes are expanded in a copy of the rule in accordance with the actual match and its application is attempted (lines 15–24 in Fig. 13). If such an application succeeds, the result is given as result of the application of the rule, that is, it is used to replace the current system (line 7). If the application in the condition fails, the rule fails. To understand this rule, we need to introduce some additional Maude machinery and some auxiliary functions.

First, in addition to equality checks, Maude rule conditions may include so-called matching equations using the operator :=. Given a pattern term P (a canonical term possibly with free variables) and a term T which may use variables in the left-hand side of the rule and also variables introduced in previous matching conditions, the condition expression P := T evaluates the term T and tries to match its result to the pattern P. If P is a variable, it works like a *let* or *where* clause to assign that value to the variable so that it can later be used. If a more general pattern is used, the match may result in the simultaneous assignment of values to multiple variables. In the rule in Fig. 13, we can see how matching conditions are used several times. First, it is used to refer to the left-hand side of the rules as { Conf′ } (in lines 9-13), then to refer to the match of the rule as Subst, and finally to get the result of the application of the unfolded rule (line 15).

In Maude, terms and modules have a metarepresentation that we can manipulate as regular terms. Up and down functions allow us to move terms and modules between levels. For instance, given the module GENERIC-PN in which these rules are defined, the expression upModule('GENERIC-PN, false) gives us its metarepresentation. Similarly, upTerm({ Conf′ }) gives us the metarepresentation of the term { Conf′ }, and downTerm(T, { none }) moves down the result of the application of the unfolded rule T. The built-in function metaApply(M,T,L,S,N) returns the Nth solution of applying rule L in module M on term T using the substitution S to constraint the application of the rule. Then, assuming that the makeModule function takes the module with the rules and expands the boxes of the indicated rule as required, the call to metaApply in the last matching condition in the rule in Fig. 13 will apply the expanded rule on the current state of the system using the original substitution, that is, forcing the same match. The metaApply function gives a triple {T, Ty, Subst} as result, where T is the term resulting from the application of the rule, Ty its type, and Subst is the complete substitution used in the application.

The FireRegularArcsBoxes rule is shown in Fig. 14. It is similar to the rule without boxes explained in Sect. 3.3, but now it also includes information on the boxes in the FROM (lines 5–6) and TO (lines 13–14) parts. The boxes representation of the FROM part includes references to objects in the level(L3) Model object in the LHS, whist the boxes rep-

resentation of the TO part has references to objects in the corresponding model object in the RHS.

Boxes are specified as a collection of terms of the form box[C]{OS}, with C being the cardinality of the box and OS the set of identifiers of the objects (nodes and relations) and nested boxes within the box. The makeModule function is a metalevel function that operates on the metarepresented module, expanding the indicated rule by unfolding the boxes in it. It proceeds recursively, removing one box level at a time. As we have seen above, the cardinality of a box specifies the number of replicas of that box that we need to generate. After a box is expanded, the cardinalities of the next-level boxes can be evaluated. The operation is repeated until no further boxes are left. Once all boxes in a rule are completely expanded, the application of the rule is attempted in one single step.

Notice that boxes in right-hand sides may, and in fact do in the FireRegularArcsBoxes rule in Fig. 14, contain identifiers of new objects. Notice also that the counter object is updated according to the number of objects being created, either inside or outside boxes.

4 Execution and analysis of a Petri net model

The capability to execute MultEcore systems and use the powerful tools that Maude implements for reachability and model checking of the multilevel models takes our infrastructure to a next step. Furthermore, modelling behavioural languages, such as Petri nets, becomes interesting if one can transfer simulation and analysis onto the concrete system models. This section shows how to execute and analize such models. In particular, we demonstrate these capabilities on a case study, namely, a gas station model adapted from [31]. We consider the complete modelling cycle, where the modeller sketches and designs the multilevel hierarchy that represents the system, then specifies the behaviour by means of transformation rules, and then automatically transforms its setting to Maude where simulation and execution can be done to verify and analyse the system.

The concrete syntax of the model is depicted in Fig. 15. This model would be located at level 4 of the PN hierarchy, as an instance of the *reset-inhibitor-petri-nets* model (Figs. 6, 18c). The PN model represents a system in which car tanks get filled up at a gas station. The station has a tank with a maximum capacity, which can be evacuated for cleaning operations and then replenished. If there is no car in the station, a new car can arrive and set its indicator on. Once the car's tank is filled, it leaves the station.

The initial marking of the model, depicted in Fig. 15, has 4 tokens in the place Station Tank, which is its full capacity. For a car tank to be refuelled, there must be a token in the Fuel Indicator On place. This can only happen if the transition Turn Fuel Indicator On has been fired, and for this

```
 1 crl [FireRegularArcs] :
 2   { < level(L1) : Model | ... >              ---- left-hand side as for the rule without boxes
 3     < level(L2) : Model | ... >
 4     < level(L3) : Model | ... >
 5     < counter : Counter | ... >
 6     Conf }
 7 => downTerm(T, { none })
 8 if (tr . outArcs -> size() = tr . outRegularArcs -> size()) /\ (tr . inArcs -> size() = tr . inRegularArcs -> size())
 9 /\ Conf' := < level(L1) : Model | ... >
10               < level(L2) : Model | ... >
11               < level(L3) : Model | ... >
12               < counter : Counter | ... >
13               Conf
14 /\ Subst := ( ... )                          ---- match of the rule
15 /\ { T, Ty, Subst' } := metaApply(
16                           makeModule(
17                             upModule('GENERIC-PN, false),
18                             upTerm({ Conf' }),
19                             'FireRegularArcsBoxes,
20                             Subst),
21                           upTerm({ Conf' }),
22                           'FireRegularArcsBoxes,
23                           Subst,
24                           0) .
```

Fig. 13 Excerpt of the first of the Maude rewrite rule corresponding to the *Fire regular transition* MCMT rule

```
 1 rl [FireRegularArcsBoxes] :
 2   { < level(L1) : Model | ... >     ---- as the left-hand side of the rule without boxes
 3     < level(L2) : Model | ... >
 4     < level(L3) : Model | ... >
 5     boxes((tr . outRegularArcs -> size()]{ O27, O29, O39, O40 },                    ---- box information
 6         box[tr . inRegularArcs -> size()]{ O32, O34, O36, O38, box[a1c]{ O22, O29 } } ))
 7     < counter : Counter | value : N >
 8     Conf }
 9 =>
10   { < level(L1) : Model | ... >     ---- as the right-hand side of the rule without boxes
11     < level(L2) : Model | ... >
12     < level(L3) : Model | ... >
13     boxes(( box[tr . outRegularArcs -> size()]{ O27, O29, O39, O40, box[aoc]{ oid(L3, N), oid(L3, N + 2)} },
14         box[tr . inRegularArcs -> size()]{ O32, O34, O36, O38 } ))
15     < counter : Counter | value : N + 13 > Conf } .
```

Fig. 14 Excerpt of the second of the Maude rewrite rule corresponding to the *Fire regular transition* MCMT rule

to happen, there cannot be any token either in the Car Tank place (the tank is empty) or in the Fuel Indicator On place (the indicator is off). This is modelled by the two inhibitor arcs connected to the Turn Fuel Indicator On transition. Then, once the indicator has been turned on, we can only progress by firing Fuel Car transition, which makes the Car Tank filled (i.e., gets one token). Ultimately, once the car tank is full, the Leave Station transition can be fired, leading the car to exit the station by putting a token into the Outside Station place.

4.1 Execution using Maude

Given a model with an initial marking, we can simulate it. In fact, we have two ways of doing so. We can let the default strategy choose the rules to apply, and the way in which to apply them, or we can force a specific sequence of rule applications.

Rule rewriting is a highly non-deterministic process, and in general, at every step many rules may be applicable (enabled). Moreover, since a rewrite system may be non-terminating, as is the case of the gas station example, a

maximum number of rewriting steps to be taken may be specified.

A finer control on rule application may sometimes be desirable. In MultEcore, we may specify the sequence of rules to be applied, which can be selected from a list of possible ones. As we states in Sect. 3.4, for any rule with boxes, several Maude rules are generated, corresponding to the different combinations of boxes with cardinality zero. Thus, for the *Fire reset/inhibitor transition* rule in Fig. 7, seven rules are generated, e.g., for transitions with no regular arcs (FireResetInhibitorArcs), for transitions with no inhibitor and no reset arcs (FireRegularArcs), for transitions with no reset and no regular arcs (FireInhibitorArcs). Then, we can specify the strategy with which we desire to rewrite our initial marking model by indicating the corresponding sequence of rule labels. For example, we can guide the execution from an initial marking given the following sequence:

```
FireInhibitorArcs  --------- Fuel Indicator On
FireRegularArcs  ---------- 1st car fuelled (Car Tank)
FireRegularArcs  ---------- 1st car leaves (Outside Station)
FireInhibitorArcs  --------- Fuel Indicator On
```
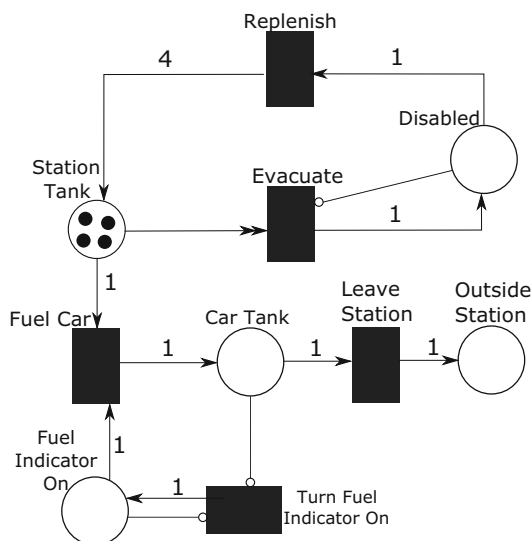
Fig. 15 Gas station model with initial marking

```
FireRegularArcs  ----------- 2nd car fuelled (Car Tank)
FireRegularArcs  ----------- 2nd car leaves (Outside Station)
FireResetInhibitorArcs ---- Disabled
FireRegularArcs  ----------- Station tank filled
```

The comments on the right indicate the corresponding effect on the model. The MultEcore integration with Maude allows us to specify sequence of rules to automatically generate the desired model state and get its graphical representation right away. A screenshot of the MultEcore tool is depicted in Fig. 16, in which we can see how rule labels are selected from a list corresponding to, the above sequence of rules. Note that the MultEcore syntax of the initial state (depicted in Petri nets syntax in Fig. 15) is in the background of Fig. 16. Pressing Finish on such a wizard automatically provides the model state in the MultEcore syntax, shown in Fig. 17, where we have four tokens (token19211…token19214) in the Station Tank place, representing that the station tank has been replenished (highlighted at the top-left of Fig. 17), and four tokens in the Outside Station place, representing that 4 cars have left the station (highlighted at the bottom-right of Fig. 17).

4.2 Reachability analysis

To perform a more exhaustive verification of the model, we can perform reachability analysis and bounded model-checking of invariants, with which we can check safety properties. Specifically, we can study the reachability of given states using the *search* command, where the states to check can be specified both using patterns or conditions on the states. The search command explores the reachable state space following a breadth-first strategy.

To carry out a search, we need to provide: (i) the model from which to initiate the search, (ii) the maximum depth

of the search (even for terminating systems, a search may take a long time), (iii) the pattern model to be reached (a model with variables), and (iv) an optional property that has to be satisfied by the reached state. Since the structure of the PN does not change along the execution, we can specify our pattern model leaving as variables the tokens in each place. Then, the condition to satisfy at the target state may be specified as an OCL expression. As result, MultEcore will determine whether such a state is reachable, in the specified number of steps, and if so, it may provide the specific path leading to the state found.

Let us see how we can use the search command of Maude to verify properties on the gas station example. For example, we can verify that, starting from the marking depicted in Fig. 15, the system can reach a marking where four cars have left the station and the station tank is again full to continue fuelling further cars. To do that, we just need to select the initial model, the target pattern model, and write, for example, the following OCL boolean expression:

```
id(4, "Station Tank").tokens->size() = 4
and id(4, "Outside Station").tokens->size() = 4
```

MultEcore responds positively, and provides the sequence of rule names leading to such a solution:

```
FireInhibitorArcs
FireRegularArcs
FireRegularArcs
FireInhibitorArcs
FireRegularArcs
FireRegularArcs
FireInhibitorArcs
FireRegularArcs
FireRegularArcs
FireInhibitorArcs
FireRegularArcs
FireRegularArcs
FireResetInhibitorArcs
FireRegularArcs
```

Notice that this is the path to one of the possible states satisfying this condition, which, like in this case, may be not unique.

We can also check whether certain misbehaviours may occur. Below we list properties we have verified in order to assure the correctness of the model:

– **Property 1.** It is not possible to find a state where, simultaneously, the places Station Tank and Disabled contain tokens. This would imply that either Evacuate or Replenish transition could be fired more than once in a row, which is not the behaviour we expect for this specific system. We can check this property using the same initial and pattern markings as before together with the following property:

```
id(4, "Station Tank").tokens->size() > 0
and id(4, "Disabled").tokens->size() > 0
```
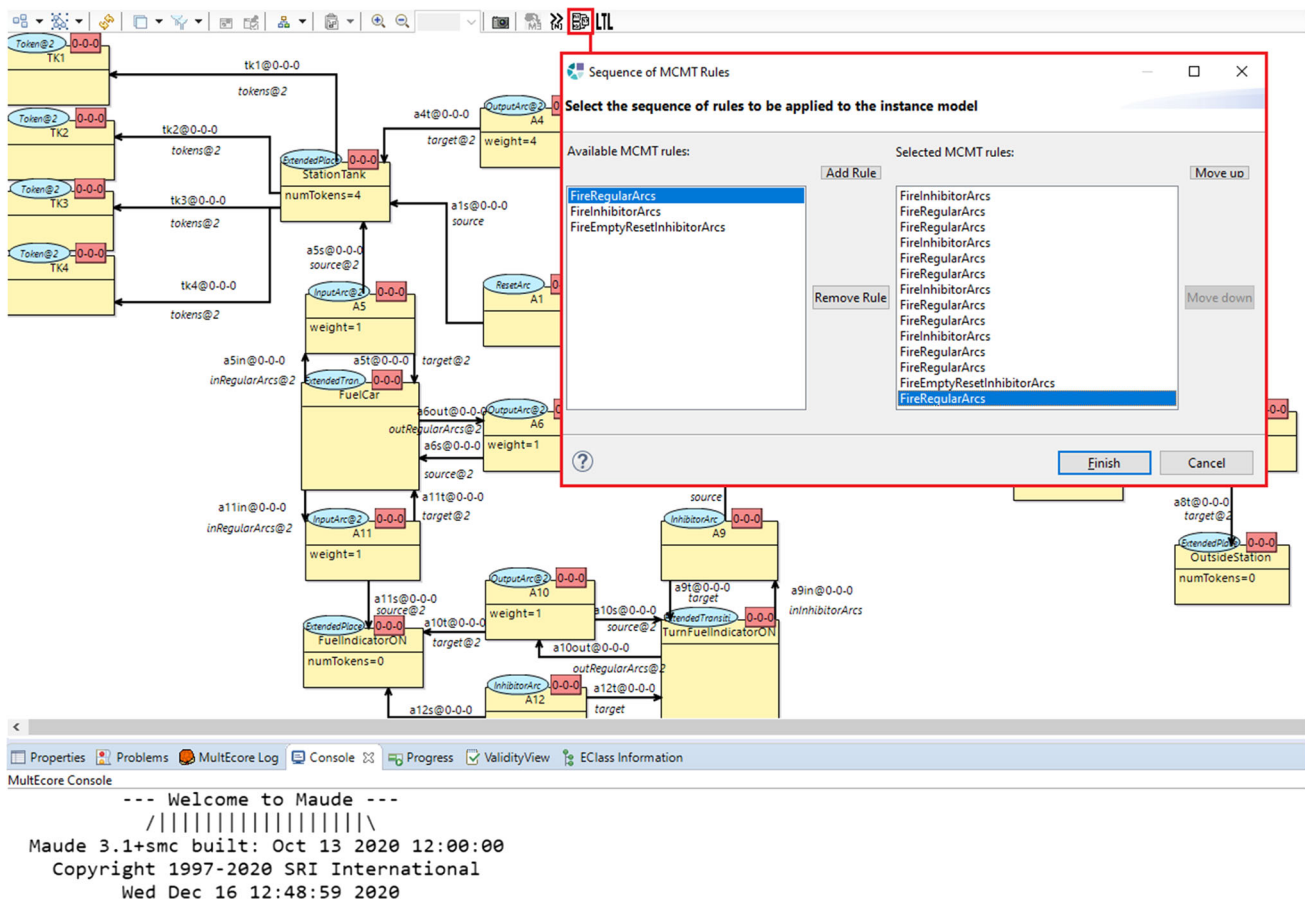
Fig. 16 MultEcore screenshot for sequence of rules-based execution

Since the system is not terminating, and we expect not to find it, then we specify a bound, e.g., of 20. The answer from MultEcore is 'false', indicating that such a state cannot be reached within the given depth.

– **Property 2.** It is not possible to find a state where either the Fuel Indicator ON, Car Tank, and Disabled places have more than one token. In other words, the indicator is either ON or OFF, the car tank is either full or not, and we cannot disable the station consecutively two times. This property can be verified as the previous one using the following OCL expression:

```
id(4, "Fuel Indicator ON").tokens->size() > 1
or id(4, "Car Tank").tokens->size() > 1
or id(4, "Disabled").tokens->size() > 1
```

Again, for a bound of 20, the answer from MultEcore is 'false'.

– **Property 3.** We can find a state where more than 5 cars have successfully exited the station.

```
id(4, "Outside Station").tokens->size() > 5
```

The answer obtained is positive, and the path to the first found solution is provided.

All the properties listed above have provided the expected results, which further validates our model.

4.3 System abstraction for unbounded analysis

In the previous section we have carried out bounded model checking of several behavioural properties. Although limited to a maximum depth, the verification of the properties checked using the search command greatly increase the confidence in the correctness of the system. However, bounded model checking is an incomplete procedure, since a counterexample could exist at a larger search depth. The problem in this case is that the state space is infinite, and therefore, we cannot complete the analysis in this way. One way to fully verify the system is using a finite abstraction of it, that is, on an appropriate quotient of the original system whose set of reachable states is finite. The method proposed in [44] creates an abstraction of the original system by adding a set of equations to collapse the infinite set of reachable states into a finite set. The specification, extended with these equations,
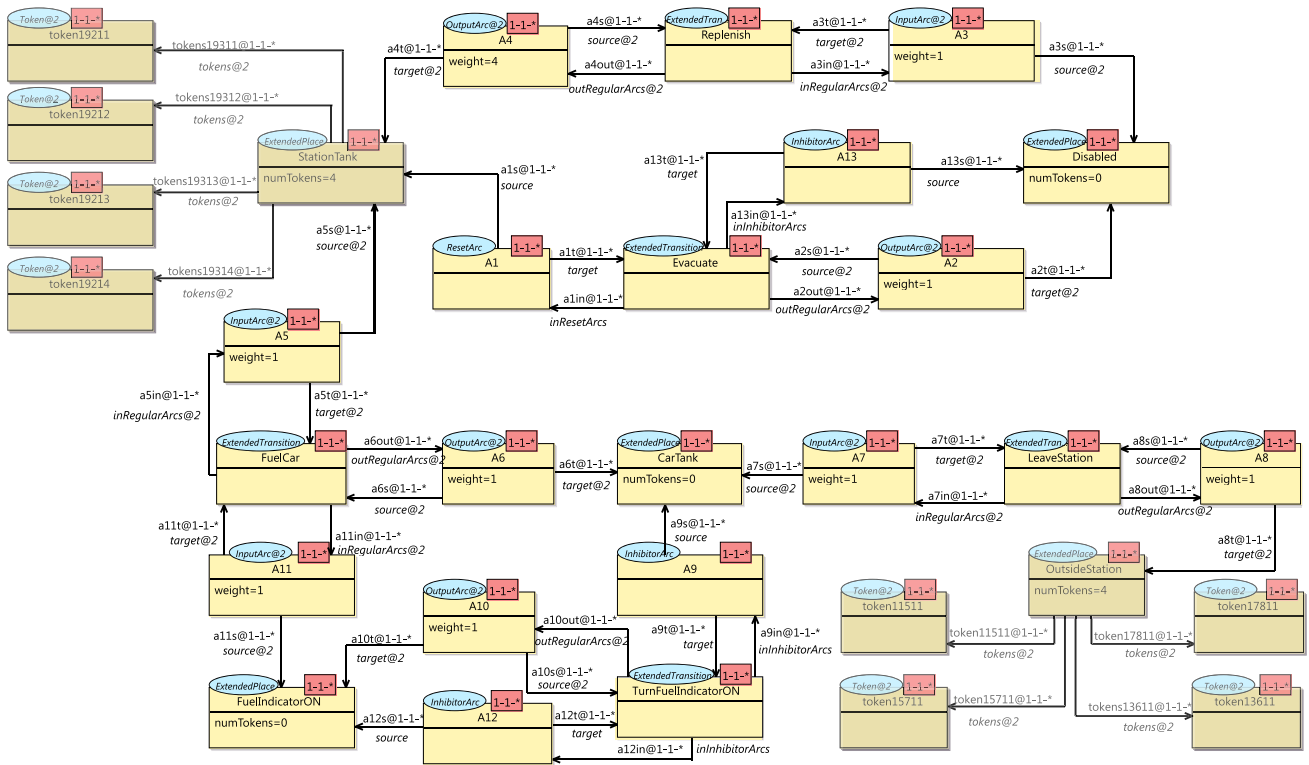
Fig. 17 MultEcore syntax of the Gas station model state after execution

need to still satisfy the usual executability conditions—the equations must be ground Church-Rosser and terminating, and the rules should be ground coherent with them—but the procedure is quite simple, and the abstraction is then correct by construction. The method is valid both for the verification of invariants and LTL formulas with an additional deadlock-freedom requirement. Indeed, an automatic procedure to complete specifications so that the requirement is satisfied has been given in [15].

The key idea about abstraction for invariant verification is that if we can verify an invariant on the abstracted specification—the specification with the equations defining the abstraction—then it also holds in the original specification. The implication, however, only works in one direction. If we find a counterexample in the abstracted system it does not necessarily mean that a counterexample exists for the original system.

Let us apply the technique to our example. There are several reasons why our Petri nets system is infinite. On the one hand, we have that tokens get accumulated in the Outside Station place, since every time a car gets its tank filled, the car leaves the station and a new token is added to the place. On the other hand, since we are representing tokens as objects, every time a new token is created it gets a new unique identifier. Thus, even though from the Petri nets point of view, tokens are anonymous dots in a place, in our representation

tokens are objects that have names and identifiers, and the counter object keeps getting its value attribute increased. We may, however, abstract from this information, since neither the tokens' names nor identifiers are relevant, nor are we really concerned about the number of tokens we have in the Outside Station place. For the operation of the Petri net, the number of tokens in the other places is not relevant either. Specifically, since all arc weights are one, the analysis would be the same if having four or three tokens in the Station Tank place.

We introduce equations that abstract the Petri net system is the following way: (1) The Outside Station place gets its number of tokens decremented if it is bigger than one, (2) the number of tokens in the Station Tank place becomes 3 if it gets 4 tokens, (3) names and identifiers of token objects are reset into a range of values not used by the counter object, and (4) the counter object gets its value attribute restarted to its initial value. Notice that when we eliminate or rename a token, we also act on the relation object associating it to the place in which it is located. In this way, we not only make all places to have either zero or one tokens in them, but names and identifiers are reused from a small set of possible values.

This abstraction makes the state space finite. And we can use it to verify LTL formulas on the abstracted model. Temporal logic allows the specification of safety properties (something bad never happens) and liveness properties

(something good eventually happens), which are related to the infinite behaviour of a system. However, we need a few additional definitions first.

Kripke structures are the natural models for propositional temporal logic. We need to understand how a Kripke structure is associated to the rewrite theory specified by a Maude system module. Basically, a Kripke structure is a (total) transition system to which we have added a collection of unary state predicates on its set of states. Therefore, since the models of rewriting logic are also transition systems, we need to make explicit the type of each of the states (System in our specification) and the atomic propositions on which we define our state predicates. In our case, we use OCL boolean expressions as basic propositions, associating to each state those boolean expressions that are satisfied in such a state. For example, we may check the following properties

- **Property 1.** To check the property stating that it is always ([]) true that eventually (<>) the system gets to a state in which there is a token on either the Disabled or the Outside Station place can be checked using the following LTL formula:

  ```
  [] <> (id(4, "Disabled").tokens->size() > 0
      ∨ id(4, "Outside Station").tokens->size() > 0)
  ```

  In this case, the answer obtained is positive.

- **Property 2.** The following LTL formula states that if we reach a state in which there is a token in the Fuel Indicator On place, then eventually a state in which there is a token in the Outside Station place is reached.

  ```
  [] (id(4, "Fuel Indicator ON").tokens->size() > 0
    -> []<> id(4, "OutsideStation").tokens->size() > 0)
  ```

  In this case the response is negative. Indeed, it may happen that the Petri net loops in the upper part, evacuating and replenishing the station tank over and over again. As usual, if the formula is not true, the model checker gives a counterexample, in the form of a sequence of states.

The interested reader can find the complete outputs in [55], together with the source files used to reproduce the execution.

5 Related work

In [56], the initial steps towards the use of rewriting logic for providing a semantics of the MultEcore language which could be used for model simulation was presented. The present paper significantly extends it in several directions, but mainly in its expressivity, support for analysis, and usability. In summary, the main contributions of this paper in comparison to [56], are:

- MultEcore has now full support for the definition and manipulation of attributes. Support for the Object Constraint Language (OCL) [13] has been added for the manipulation of attribute values and for the specification of conditions, which greatly improves the expressiveness of the tool.
- The expressivity of MCMTs have been extended and improved, with respect to those used in [41,56] in the following three main ways: (i) attribute definition and manipulation, which brings additional expressivity to the specification of behaviour; (ii) rule conditions that add extra requirements for a rule to be applied; and (iii) nested boxes to handle submodel collections, improving expressiveness and reducing the proliferation of rules. A preliminary and very limited version of these boxes was presented in [56,58]. In this paper we presented a fully-operational full-fledged version of them, where boxes may appear in both sides of the rules, boxes may be nested, and each of them may have an explicit cardinality specified.
- We present in this paper a rewriting logic semantics of MLM hierarchies and MCMT rules through their representation in Maude. This formal representation of MultEcore models allows us to execute and analyse such models using Maude's formal tools. Specifically, we have illustrated the application of model checking, both in the search for states satisfying provided conditions, and for the verification of linear temporal logic formulas.
- A bidirectional transformation between MultEcore MLM models and Maude specifications has been developed. This functional infrastructure connects MultEcore to Maude, allowing us, not only to design our multilevel modelling hierarchy and specify its MCMTs, but also to simulate the specified systems and analyse and verify them using several techniques. Within our infrastructure, we encapsulate Maude as a background process that handles the instructions and return the execution and analysis results, given by the interface that the user uses to interact. Whist there was some basic infrastructure for this in [56], this is now a mature tool, not only more efficient and configurable, but covering all the features of the language.

Even though there exist a plethora of MLM approaches and tools, only a few of them support DSML behaviour specification and execution. Melanee [2] is one of the most advanced tools for MLM. The tool supports a variant of OCL with deep semantics (DeepOCL) which has been integrated with the Atlas Transformation Language (ATL) for model transformations. Lange shows in [36] how this tool can be used to check constraints spanning multiple classification levels which can be defined and executed. Although Melanee itself is not natively supporting tools for simulation/execu-

tion through the specification of the execution semantics, i.e., (multilevel) transformation rules, there are some works on top of it that aims to achieve this (see, e.g., [3]). In that work, the model execution mechanism is based on a service API and a plug-in mechanism, and the communication between the modelling and the execution environments is realised using socket-based communication. We provide in our approach the whole set of tools necessary to directly be able to define the structure of the multilevel hierarchy, specify the multilevel model transformation rules (MCMTs), execute/simulate the models, and analyse the system.

The MetaDepth tool [17] is a well-known framework within the MLM community. It is integrated with the Epsilon languages [25], which permits using the Epsilon Object Language (EOL) as an action language to define behaviour for metamodels, as well as the Epsilon Validation Language (EVL) for expressing constraints. Both EOL and EVL are extensions of OCL. The approach implements the interface of the connectivity layer in a way to make EOL aware of the multiple ontological levels providing it with a multilevel nature. However, the authors of [17] state that MetaDepth can be used as a normal two-level meta-modelling environment when it comes to the execution of behaviour of the models. Thus, for the actual execution they would have to flatten their multilevel language to a two-level version in order to run the models. To the best of our knowledge there is not yet a MLM tool that supports or integrates model checking and analysis capabilities within its MLM tool-set.

Other authors have attempted to handle pattern identification and specification to define reusable model transformation rules. There exist a diverse set of approaches that bring solutions to pattern definition and application in the context of graph transformations. In [29], Guerra and de Lara explore *recursion* as a graph transformation mechanism. They provide double pushout (DPO) rules with base and recursive conditions, together with mechanisms to pass the matching between successive recursion steps. Lindqvist et al. [37] propose the star operator, which is suited to find repetitive occurrences of a specific modelling pattern. However, the star operator is only defined for matching model extracts, and not to perform transformations. In [28], Grønmo, Krogdahl, and Møller-Pedersen present a collection operator for graph transformation and show its usage to a variety of Coloured Petri nets. Using this operator, it is possible to match several similar structures within the model. They theoretically define how nesting would work by producing an ad-hoc rule that would fit to the specific case. We follow a similar approach by defining the rule in a generic way and then the transformation engine provides also a generic version in the Maude specification. The advantage of our approach with respect to the collection operator is that we do not physically produce an unfolded rule, but it is dynamically unfolded and used at run-time. It is during the matching at run-time when the rule is unfolded guided by the cardinalities provided in the rule. In [51], Rensink and Kuperus propose a transformation language that uses an amalgamation scheme for nested graph transformation rules, where pattern elements are combined with universal and existing quantifiers. The transformation language is used in the GROOVE tool. Henshin [61] is an inplace model transformation language for the EMF. Among other features, it implements a *rule-nesting* mechanism [1] that provides a for-each operator for rules. In nested rules, the outer rule is referred to as *kernel rule* and the inner rule as *multi rule*. During execution of a nested rule, the kernel rule is matched and executed once. Afterwards, the match is used as a starting point to match the multi-rule as often as possible and execute it for each match.

There are several traditional MDSE approaches that deal with execution and verification related to our proposal. In [53], Rivera et al. use Maude to represent 2-level models to be able to simulate and perform formal analysis and model checking on them. Such work served us as inspiration and starting point. We were considering either implementing ourselves an execution engine within MultEcore or using an existing tool where to rely on. Studying some works where Maude was used and analysing how the language could be customised together with its plethora of existing capabilities made us to follow such a path.

We also analysed other mature tools in the context of model execution via operational semantics, e.g., Henshin [61] or the GEMOC Studio [10,16], which helped us to understand how the user could interact with the execution tools from an Eclipse-based application. In the context of verification, GROOVE (GRaph-based Object-Oriented VErification) [27,49] is a tool for software model checking of object-oriented systems. It can be used for modelling, analysis and verification and integrates all these functionalities in an easy to use interactive GUI. While we already integrate into MultEcore some Maude functionalities for execution and verification, we still have to work in this direction to ease the process to the modeller. We see GROOVE as a viable approach to achieve a better usability degree in MultEcore.

6 Conclusions and future work

In this paper, we have presented an infrastructure for execution and analysis of multilevel modelling languages.

To make this possible, we have integrated Maude into the MultEcore tool, allowing us not only to define multilevel hierarchies and their behaviours by means of MCMTs, but also carry out simulation and model checking using different analysis techniques. An important element of our approach is that Maude is used as a backend tool, hidden to the user such that the interaction is entirely done with MultEcore. This makes the modeller unaware of the Maude details. Although

in the last years several traditional two-level tools have provided support for the whole cycle of behavioural DMSLs (from design to simulation and verification), to the best of our knowledge this is the first work where a MLM tool incorporates capabilities to perform model checking and other formal analyses. We find that the work presented in this paper can open new doors to the MLM field, as this tool can be used to define behavioural multilevel DSMLs, execute them, and verify them.

Apart from the major improvements in the infrastructure itself, we have improved and extended the MCMTs capabilities, by allowing nested boxes to represent collections, incorporated attribute manipulations, and specification of conditions. Basic support for OCL is also provided, which is very useful for the manipulation of attributes, the specification of box cardinalities, the specification of rule conditions, and the specification of expressions and conditions to be used in the formal checks, including LTL formulas.

To validate and demonstrate that our infrastructure works and that actual execution and analysis can be carried out, we have provided a case study where a Petri net model that captures a gas station is simulated applying consecutively the MCMT rules defined in MultEcore. The goal of this case study has been to evaluate the usability and practicability of the developed infrastructure. We are already considering how to evaluate our tool against other MLM approaches that allow the modeller to perform execution on models by defining in-place model transformations. We have validated and verified the modelled system using reachability analysis and model checking techniques on an abstracted version of the model. We refer the reader to the main MultEcore webpage [54] for further details and examples.

We plan to integrate into MCMTs the capabilities that a programming language brings such as reasoning about functions, expressions, type specifications, and data manipulations. Our current implementation of certain OCL functions represents a step towards this goal. We are already working on adapting the complete mOdCL (Maude + OCL) [24] to our multilevel infrastructure so that we can make use of the full power of OCL in a Multilevel context. This would allow us, for instance, to specify Coloured Petri nets [32] which combine classical Petri nets with a programming language [63].

In the PN example presented in this paper, the behaviour of regular PNs is specified by one single MCMT, and the behaviour of reset/inhibitor PNs by another one. However, this is an exceptional case, since the behaviour of systems is typically described by a bigger number of MCMT rules, each one specifying one of the possible actions in the DSML being modelled. In systems like the robot system modelled in [40], the production line system modelled in [56], etc., one has many MCMTs, defined at different levels of abstraction. In these cases, the possibilities for horizontal/vertical extensions was clearly illustrated, showing how MCMTs defined for models in higher levels are applicable on lower levels. In the case of PNs, the only action is the firing of transitions, and the fact that the types of incoming arcs may be arbitrarily mixed, force us to model the action in one single MCMT per type of PN. However, the PNs case points an interesting issue, which we believe is worth investigating in the future. The fact that the *Fire reset/inhibitor transition* rule (Fig. 7) extends the the *Fire regular transition* MCMT (Fig. 4) makes us think on the possibility of developing mechanisms for MCMT extension, which may further improve the degree of modularity and reusability of MCMTs.

Ultimately, we plan to further advance on the interface that connects MultEcore to Maude, bringing more advanced functionalities such as an interactive editor, a smoother experience to the user with the graphical editor and additional Maude capabilities to, for instance, customise the strategy language and have more control on the execution and analysis.
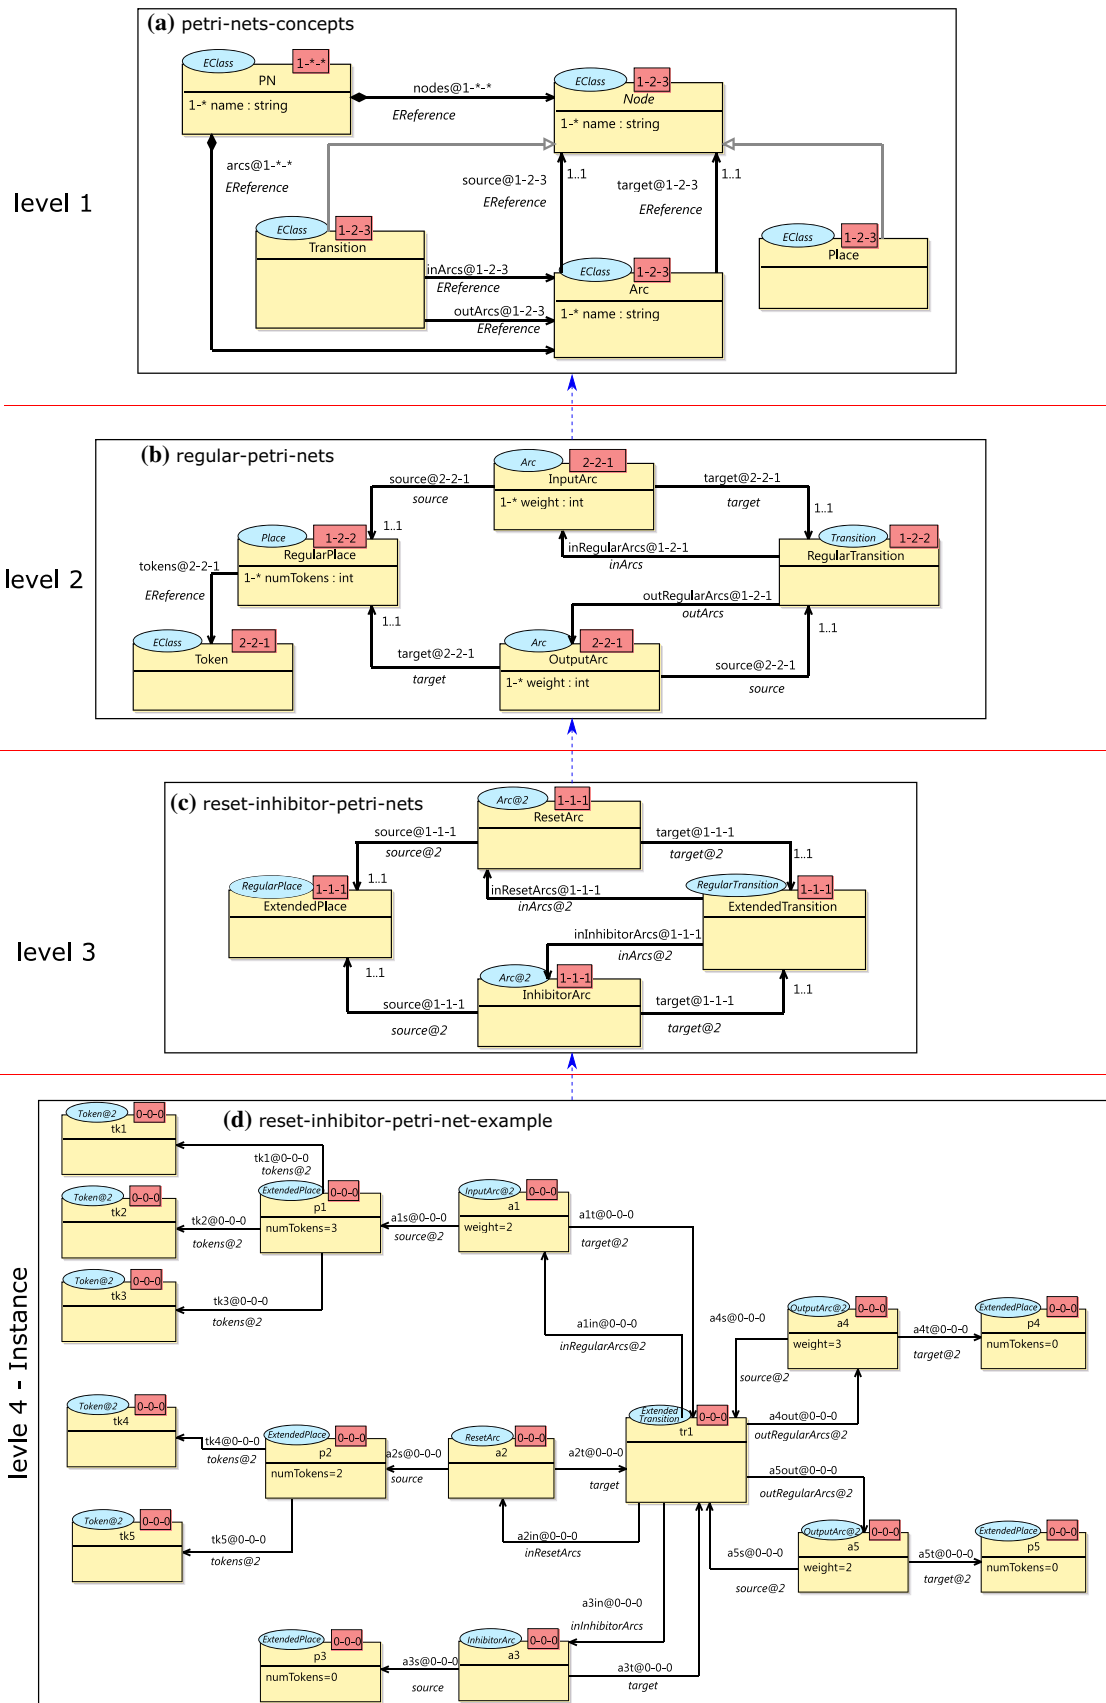
A Petri nets multilevel hierarchy

See Fig. 18.

Fig. 18 Petri nets multilevel hierarchy

References

1. Arendt, T., Biermann, E., Jurack, S., Krause, C., Taentzer, G.: Henshin: advanced concepts and tools for in-place EMF model transformations. In: 13th International Conference on Model Driven Engineering Languages and Systems (MODELS 2010), pp. 121–135 (2010). https://doi.org/10.1007/978-3-642-16145-2_9

2. Atkinson, C., Gerbig, R.: Flexible deep modeling with Melanee. In: Betz, S., Reimer, U. (eds.) Modellierung 2016, LNI, vol. 255, Gesellschaft für Informatik, Bonn, pp. 117–122 (2016)

3. Atkinson, C., Gerbig, R., Metzger, N.: On the execution of deep models. In: 1st International Workshop on Executable Modeling co-located with ACM/IEEE 18th International Conference on Model Driven Engineering Languages and Systems (MODELS 2015), pp. 28–33 (2015)

4. Atkinson, C., Kühne, T.: Processes and products in a multi-level metamodeling architecture. Int. J. Software Eng. Knowl. Eng. **11**(06), 761–783 (2001)

5. Atkinson, C., Kühne, T.: The essence of multilevel metamodeling. In: UML 2001—The Unified Modeling Language, Modeling Languages, Concepts, and Tools, pp. 19–33 (2001). https://doi.org/10.1007/3-540-45441-1_3

6. Atkinson, C., Kühne, T.: Reducing accidental complexity in domain models. Softw. Syst. Model. **7**(3), 345–359 (2008)

7. Atkinson, C., Kühne, T.: In defence of deep modelling. Inf. Softw. Technol. **64**, 36–51 (2015). https://doi.org/10.1016/j.infsof.2015.03.010

8. Atkinson, C., Kühne, T.: On evaluating multi-level modeling. In: Proceedings of MULTI @ MODELS, pp. 274–277 (2017)

9. Bernardinello, L., de Cindio, F.: A survey of basic net models and modular net classes. In: Advances in Petri Nets 1992, The DEMON Project, Springer, pp. 304–351 (1992). https://doi.org/10.1007/3-540-55610-9_177

10. Bousse, E., Wimmer, M.: Domain-level observation and control for compiled executable dsls. In: Kessentini, M., Yue, T., Pretschner, A., Voss, S., Burgueño L. (eds.), 22nd ACM/IEEE International Conference on Model Driven Engineering Languages and Systems, MODELS 2019, Munich, Germany, September 15–20, 2019, IEEE, pp. 150–160 (2019). https://doi.org/10.1109/MODELS.2019.0006

11. Brambilla, M., Cabot, J., Wimmer, M.: Model-Driven Software Engineering in Practice. Synthesis Lectures on Software Engineering. Morgan & Claypool Publishers (2012). https://doi.org/10.2200/S00441ED1V01Y201208SWE001

12. Cabot, J., Gogolla, M.: Object constraint language (OCL): a definitive guide. In: Bernardo, M., Cortellessa, V., Pierantonio, A. (eds.), Formal Methods for Model-Driven Engineering—12th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2012, Bertinoro, Italy, June 18–23, 2012. Advanced Lectures, Lecture Notes in Computer Science, vol. 7320, Springer, pp. 58–90 (2012). https://doi.org/10.1007/978-3-642-30982-3_3

13. Clark, T., Warmer, J.: Object Modeling With the OCL: The Rationale Behind the Object Constraint Language, vol. 2263, Springer (2003)

14. Clavel, M., Durán, F., Eker, S., Lincoln, P., Martı-Oliet, N., Meseguer, J., Quesada, J.F.: Maude: specification and programming in rewriting logic. Theoret. Comput. Sci. **285**(2), 187–243 (2002)

15. Clavel, M., Durán, F., Eker, S., Lincoln, P., Martí-Oliet, N., Meseguer, J., Talcott, C.L. (eds.) All About Maude—A High-Performance Logical Framework, How to Specify, Program and Verify Systems in Rewriting Logic. Lecture Notes in Computer Science, vol. 4350, Springer (2007). https://doi.org/10.1007/978-3-540-71999-1

16. Combemale, B., Barais, O., Wortmann, A.: Language engineering with the GEMOC studio. In: 2017 IEEE International Conference on Software Architecture Workshops, ICSA Workshops 2017, Gothenburg, Sweden, April 5–7, 2017, IEEE Computer Society, pp. 189–191 (2017). https://doi.org/10.1109/ICSAW.2017.61

17. de Lara, J., Guerra, E.: Deep meta-modelling with MetaDepth. In: Objects, Models, Components, Patterns, lncs, vol. 6141, Springer, pp. 1–20 (2010). https://doi.org/10.1007/978-3-642-13953-6_1

18. de Lara, J., Guerra, E.: Generic meta-modelling with concepts, templates and Mixin layers. In: Model Driven Engineering Languages and Systems—13th International Conference, MODELS, pp. 16–30 (2010). https://doi.org/10.1007/978-3-642-16145-2_2

19. de Lara, J., Guerra, E., Cuadrado, J.S.: When and how to use multilevel modelling. ACM Trans. Softw. Eng. Methodol. (TOSEM) **24**(2), 12 (2014)

20. de Lara, J., Vangheluwe, H.: AToM 3: a tool for multi-formalism and meta-modelling. In: International Conference on Fundamental Approaches to Software Engineering, Springer, pp. 174–188 (2002)

21. Durán, F., Eker, S., Escobar, S., Martí-Oliet, N., Meseguer, J., Rubio, R., Talcott, C.L.: Programming and symbolic computation in Maude. J. Log. Algebraic Methods Program. (2020). https://doi.org/10.1016/j.jlamp.2019.100497

22. Durán, F., Garavel, H.: The rewrite engines competitions: a rectrospective. In: Beyer, D., Huisman, M., Kordon, F., Steffen, B. (eds.) Tools and Algorithms for the Construction and Analysis of Systems—25 Years of TACAS: TOOLympics, Held as Part of ETAPS 2019, Proceedings, Part III, Lecture Notes in Computer Science, vol. 11429, Springer, pp. 93–100 (2019). https://doi.org/10.1007/978-3-030-17502-3_6

23. Durán, F., Rocha, C., Álvarez, J.M.: Tool interoperability in the maude formal environment. In: Corradini, A., Klin, B., Cîrstea, C. (eds.) Algebra and Coalgebra in Computer Science—4th International Conference, CALCO 2011, Winchester, UK, August 30–September 2, 2011. Proceedings, Lecture Notes in Computer Science, vol. 6859, Springer, pp. 400–406 (2011). https://doi.org/10.1007/978-3-642-22944-2_30

24. Durán, F., Roldán, M.: Validating OCL constraints on Maude prototypes of UML models. Universidad de Málaga, Technical Report (2012)

25. The Epsilon Object Language (EOL). https://www.eclipse.org/epsilon/doc/eol/

26. Garavel, H., Tabikh, M., Arrada, I.: Benchmarking implementations of term rewriting and pattern matching in algebraic, functional, and object-oriented languages—the 4th rewrite engines competition. In: Rusu, V. (ed.) Rewriting Logic and Its Applications—12th International Workshop, WRLA 2018, Held as a Satellite Event of ETAPS, Proceedings, Lecture Notes in Computer Science, vol. 11152, Springer, pp. 1–25 (2018). https://doi.org/10.1007/978-3-319-99840-4_1

27. Ghamarian, A.H., de Mol, M., Rensink, A., Zambon, E., Zimakova, M.: Modelling and analysis using GROOVE. Int. J. Softw. Tools Technol. Transf. **14**(1), 15–40 (2012). https://doi.org/10.1007/s10009-011-0186-x

28. Grønmo, R., Krogdahl, S., Møller-Pedersen, B.: A collection operator for graph transformation. Softw. Syst. Model. **12**(1), 121–144 (2013). https://doi.org/10.1007/s10270-011-0190-3

29. Guerra, E., de Lara, J.: Adding recursion to graph transformation. ECEASST (2007). https://doi.org/10.14279/tuj.eceasst.6.56

30. Halder, A., Venkateswarlu, A.: A study of petri nets modeling analysis and simulation. Department of Aerospace Engineering Indian Institute of Technology Kharagpur, India (2006)

31. Hee, van, K., Leurs, M., Post, R.: Yasper : Yet another smart process editor (poster). In: 2005 Symposium on Verification and validation of software systems (VVSS 2005) (2005)

32. Jensen, K., Kristensen, L.M., Wells, L.: Coloured petri nets and CPN tools for modelling and validation of concurrent systems. Int. J. Softw. Tools Technol. Transf. **9**(3), 213–254 (2007). https://doi.org/10.1007/s10009-007-0038-x

33. Kelly, S., Tolvanen, J.: Domain-Specific Modeling: Enabling Full Code Generation. Wiley, New York (2008)

34. Kühne, T.: Exploring potency. In: Proceedings of the 21th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems, MODELS, pp. 2–12 (2018). https://doi.org/10.1145/3239372.3239411

35. Kühne, T.: A story of levels. In: Proceedings of MULTI @ MODELS, pp. 673–682 (2018)

36. Lange, A.: dACL: the deep constraint and action language for static and dynamic semantic definition in Melanee (2016). https://madoc.bib.uni-mannheim.de/43490/. Unpublished

37. Lindqvist, J., Lundkvist, T., Porres, I.: A query language with the star operator. ECEASST (2007). https://doi.org/10.14279/tuj.eceasst.6.55

38. Macías, F.: Multilevel modelling and domain-specific languages. PhD Thesis, Western Norway University of Applied Sciences and University of Oslo (2019)

39. Macías, F., Rutle, A., Stolz, V.: Multilevel modelling with MultEcore: a contribution to the MULTI 2017 challenge. In: Proceedings of MULTI @ MODELS, pp. 269–273 (2017)

40. Macías, F., Rutle, A., Stolz, V., Rodríguez-Echeverría, R., Wolter, U.: An approach to flexible multilevel modelling. Enterp. Model. Inf. Syst. Archit. Int. J. Concept. Model. **13**, 10:1–10:35 (2018). https://doi.org/10.18417/emisa.13.10

41. Macías, F., Wolter, U., Rutle, A., Durán, F., Rodriguez-Echeverria, R.: Multilevel coupled model transformations for precise and reusable definition of model behaviour. J. Log. Algebraic Methods Program. **106**, 167–195 (2019). https://doi.org/10.1016/j.jlamp.2018.12.005

42. Meseguer, J.: Conditioned rewriting logic as a united model of concurrency. Theor. Comput. Sci. **96**(1), 73–155 (1992). https://doi.org/10.1016/0304-3975(92)90182-F

43. Meseguer, J.: Twenty years of rewriting logic. J. Log. Algebr. Program. **81**(7–8), 721–781 (2012). https://doi.org/10.1016/j.jlap.2012.06.003

44. Meseguer, J., Palomino, M., Martí-Oliet, N.: Equational abstractions. Theor. Comput. Sci. **403**(2–3), 239–264 (2008). https://doi.org/10.1016/j.tcs.2008.04.040

45. Meta Object Facility (MOF) specification 2.5.1. https://www.omg.org/spec/MOF

46. Mohagheghi, P., Gilani, W., Stefanescu, A., Fernández, M.A., Nordmoen, B., Fritzsche, M.: Where does model-driven engineering help? Experiences from three industrial cases. Softw. Syst. Model. **12**(3), 619–639 (2013)

47. Murata, T.: Petri nets: properties, analysis and applications. Proc. IEEE **77**(4), 541–580 (1989)

48. Reisig, W.: Understanding Petri Nets—Modeling Techniques, Analysis Methods. Case Studies, Springer (2013). https://doi.org/10.1007/978-3-642-33278-4

49. Rensink, A.: The GROOVE simulator: a tool for state space generation. In: Pfaltz, J.L., Nagl, M., öhlen, B.B. (eds.) Applications of Graph Transformations with Industrial Relevance, Second International Workshop, AGTIVE 2003, Charlottesville, VA, USA, September 27–October 1, 2003, Lecture Notes in Computer Science, vol. 3062, Springer, pp. 479–485 (2003). https://doi.org/10.1007/978-3-540-25959-6_40

50. Rensink, A.: The GROOVE simulator: a tool for state space generation. In: International Workshop on Applications of Graph Transformations with Industrial Relevance, Springer, pp. 479–485 (2003)

51. Rensink, A., Kuperus, J.: Repotting the geraniums: on nested graph transformation rules. In: Electronic Communication of the European Association of Software Science and Technology, vol. 18 (2009). https://doi.org/10.14279/tuj.eceasst.18.260

52. Rivera, J.E., Durán, F., Vallecillo, A.: A graphical approach for modeling time-dependent behavior of DSLs. In: Visual Languages and Human-Centric Computing, 2009. VL/HCC 2009. IEEE Symposium on, IEEE, pp. 51–55 (2009)

53. Rivera, J.E., Durán, F., Vallecillo, A.: Formal specification and analysis of domain specific models using Maude. Simulation **85**(11–12), 778–792 (2009). https://doi.org/10.1177/0037549709341635

54. Rodríguez, A., Durán, F., Kristensen, L.M.: MultEcore webpage (2021). https://ict.hvl.no/multecore/

55. Rodríguez, A., Durán, F., Kristensen, L.M.: Petri nets experiment resources: MultEcore and Maude files (2021). https://bitbucket.org/phdalejandro/no.hvl.multecore.examples.sosym.petrinets

56. Rodríguez, A., Durán, F., Rutle, A., Kristensen, L.M.: Executing multilevel domain-specific models in maude. J. Object Technol. **18**(2), 1–21 (2019). https://doi.org/10.5381/jot.2019.18.2.a4

57. Rodríguez, A., Macías, F.: Multilevel modelling with MultEcore: a contribution to the MULTI process challenge. In: Proceedings of MULTI @ MODELS, pp. 152–163 (2019). https://doi.org/10.1109/MODELS-C.2019.00026

58. Rodríguez, A., Rutle, A., Kristensen, L.M., Durán, F.: A foundation for the composition of multilevel domain-specific languages. In: MULTI@ MoDELS, pp. 88–97 (2019). https://doi.org/10.1109/MODELS-C.2019.00018

59. Roldán, M., Durán, F.: Dynamic validation of OCL constraints with mOdCL. Electron. Commun. Eur. Assoc. Softw. Sci. Technol. (2011). https://doi.org/10.14279/tuj.eceasst.44.625

60. Steinberg, D., Budinsky, F., Merks, E., Paternostro, M.: EMF: Eclipse Modeling Framework, Pearson Education (2008)

61. Strüber, D., Born, K., Gill, K.D., Groner, R., Kehrer, T., Ohrndorf, M., Tichy, M.: Henshin: A usability-focused framework for EMF model transformation development. In: 10th International Conference, ICGT 2017, pp. 196–208 (2017). https://doi.org/10.1007/978-3-319-61470-0_12

62. Syriani, E., Vangheluwe, H., Mannadiar, R., Hansen, C., Van Mierlo, S., Ergin, H.: AToMPM: a web-based modeling environment. In: MODELS-JP 2013, CEUR Workshop Proceedings, vol. 1115, pp. 21–25 (2013)

63. Ullman, J.D.: Elements of ML Programming, Prentice-Hall, Inc. (1994)

64. The Unified Modelling Language (UML) specification 2.5.1. https://www.omg.org/spec/UML

65. Van Mierlo, S., Barroca, B., Vangheluwe, H., Syriani, E., Kühne, T.: Multi-level modelling in the Modelverse. In: MULTI@ MoDELS, CEUR Workshop Proceedings, vol. 1286, pp. 83–92 (2014)

66. Verbeek, H.M.W., Wynn, M.T., van der Aalst, W.M.P., ter Hofstede, A.H.M.: Reduction rules for reset/inhibitor nets. J. Comput. Syst. Sci. **76**(2), 125–143 (2010). https://doi.org/10.1016/j.jcss.2009.06.003

67. Warmer, J., Kleppe, A.: The Object Constraint Language Second Edition: Getting Your Models Ready for MDA, Addison-Wesley Educational Publishers (2003)

68. Wolter, U., Macías, F., Rutle, A.: The Category of Typing Chains as a Foundation of Multilevel Typed Model Transformations. Technical Report 2019-417, University of Bergen, Department of Informatics (2019)

Alejandro Rodríguez is a Ph.D. student at the Western Norway University of Applied Sciences. He is currently researching in Model-driven software engineering, Multilevel Modelling and coloured Petri net fields. He is part of the Software Engineering, Sensor Networks and Engineering Computing department.

Lars Michael Kristensen received the PhD in computer science from University of Aarhus, and is currently professor in software engineering at Western Norway University of Applied Sciences. He has published more than 70 papers in strictly referred journal and conferences, is member of the Editorial Board of the TopNoC Springer journal, and is a member of the steering committee for the International Petri Nets conference. He is co-author of the most recent textbook on Coloured Petri Net and CPN Tools which is one of the most widely used software tools for modelling and validation of concurrent systems.

Francisco Durán is Full Professor at the Department of Computer Science of the University of Málaga, Spain. He received his Ph.D. degree in Computer Science from the University of Málaga in 1999, after several years as an International Fellow at SRI International, CA. He is one of the developers of the Maude system, and his research interests deal with the application of formal methods to software engineering, including topics such as cloud systems, model-driven engineering, component-based software development, open distributed programming, reflection and meta-programming, and software composition.