

# Towards Model Checking Executable UML Specifications in mCRL2

Helle Hvid Hansen\*, Jeroen Ketema<sup>†</sup>, Bas Luttik\*, MohammadReza Mousavi\* and Jaco van de Pol<sup>†</sup>  
\*Eindhoven University of Technology  
Eindhoven, The Netherlands  
{h.h.hansen,s.p.luttik,m.r.mousavi}@tue.nl  
<sup>†</sup>University of Twente  
Enschede, The Netherlands  
{j.ketema,j.c.vandepol}@ewi.utwente.nl

**Abstract**—We describe a formalisation of a subset of Executable UML (xUML) in the process algebraic specification language mCRL2. This formalisation includes class diagrams with class generalisations and state machines with send and change events. The choice of these xUML constructs is dictated by their use in the modelling of railway interlocking systems.

The long term goal is to verify safety properties of interlockings modelled in xUML using the mCRL2 and LTSmin toolsets. Initial verification of an interlocking toy example has demonstrated that the safety properties of model instances depend crucially on the run-to-completion assumptions made.

**Keywords**—Software verification and validation, Specification languages

## I. INTRODUCTION

We formalise a subset of Executable UML (xUML) [1] in the formal specification language mCRL2 [2] with the purpose of verifying safety properties. The xUML constructs covered include class diagrams with class generalisations and object associations, and state machines which consist of composite and concurrent states and have signal and change events. The mCRL2 language extends the process algebra ACP [3] with abstract data types and is accompanied by a toolset<sup>1</sup> which provides for explicit model checking, state space analysis and simulation. Symbolic model checking is provided for by the LTSmin toolset [4], [5]<sup>2</sup>.

Our work is part of INESS<sup>3</sup>, an EC FP7-funded project, which aims at developing uniform specifications of railway interlockings. One of the aims of INESS is to formally verify safety properties of generic interlockings whose functional requirements are expressed in a dialect of xUML. Several approaches to this are explored within the project. Our approach is based on model checking using mCRL2 and the current paper provides a first step in achieving this by presenting a translation from xUML to mCRL2.

Interlockings ensure that trains neither collide nor derail. They do this by preventing conflicting routes from being set over arrangements of tracks, points, etc., where a route is the basic notion used by a railway signaller to guide trains.

As in most real-world applications, the xUML models arising from the interlocking domain are of considerable size

and they express requirements that may change over time. This leads to two main concerns: (1) In order to mitigate the state space explosion problem, our translation from xUML to mCRL2 should produce specifications that are manageable by our tools while still capturing all executions allowed by the xUML model. (2) In order to make the translation from xUML to mCRL2 efficient, correct and flexible, our approach should have the potential to be automated.

Below, we focus on the first of the above concerns, although our translation has been developed with automation in mind (see Section VIII). We have tested the viability of different translations using a toy interlocking specification, which was kindly provided to us by KnowGravity Inc.<sup>4</sup> This toy specification is almost as simple as it gets, but it still shows how choices in the translation affect the size and behaviour of the model. In particular, initial results reveal that different run-to-completion assumptions give rise to a wide variety in model sizes and observable traces.

The rest of this paper is organised as follows. In Section II, we describe and motivate the use of model checking in the verification of xUML interlocking models. In Section III, we introduce the xUML constructs covered by our translation, and discuss different types of run-to-completion assumptions. Our translation to mCRL2 is described in Section IV, and in Section V, we discuss some problems related to model checking translated xUML models and give some solutions to these problems. In Section VI, we report on our observations made in translating and verifying a toy xUML interlocking model. Finally, we discuss related work and conclude in Sections VII and VIII, respectively.

## II. MODEL CHECKING xUML INTERLOCKING MODELS

In model checking, a formal model represents all possible executions of the system being modelled. In our case, we obtain such a formal model by instantiating the classes and associations of the xUML model according to a particular track layout, where a track layout is a configuration of physical and logical railway elements such as tracks, points, signals and routes (see Section VI for an example). Suitable track layouts have been designed by railways. For example, ProRail<sup>5</sup>, the Dutch railway infrastructure manager, has

<sup>1</sup><http://www.mcr12.org>

<sup>2</sup><http://fmt.cs.utwente.nl/tools/ltsmin/>

<sup>3</sup><http://www.iness.eu>

<sup>4</sup><http://www.knowgravity.com>

<sup>5</sup><http://www.prorail.nl>

designed three track layouts that together are supposed to capture all features of track layouts found in the Netherlands.

Model checking is thus limited to verification of particular model instances. It is not possible to prove statements about all instances. In spite of this, model checking can still provide valuable information: The violation of a safety property in a particular model instance shows that the xUML model is not correct in general, and traces that witness this undesired behaviour can be used to improve the model. Moreover, confidence in the correctness of the xUML model can be increased by showing that several instances satisfy the requirements. Furthermore, model checking is more thorough than simulation, since exhaustive state space exploration is performed, which is generally not possible in simulation.

### III. EXECUTABLE UML

Executable UML [1] (xUML) consists of UML class diagrams, UML state machines and an action language which complies with the UML action semantics. There are several action languages in use; we refer to [1] for a—somewhat limited—overview.

The xUML models to be translated by us are expressed in KnowGravity’s Cassandra/xUML dialect [6]. We briefly describe the modelling constructs relevant to us.

#### A. Constructs

In class diagrams, we allow for class generalisations (inheritance) and associations between classes (specifying which class instantiations may reference each other). Association classes, however, may not occur, i.e., no objects may be related to instances of associations.

State machines may contain concurrent and composite states (AND- and OR-states) and initial pseudo-states. We currently do not translate history and final pseudo-states. All UML-defined transitions may occur as far as they involve the allowed (pseudo-)states. A transition is labelled with a trigger and a sequence of actions, both of which may be empty. A trigger must be a signal event or a change event.

Signals are communicated asynchronously. A signal can be sent either by an object within the system (an internal signal) or by the environment (an external signal). Each state machine is accompanied by an event pool which stores received signals until dispatched, i.e., until they are taken from the event pool by the state machine [7, Section 13.3.25].

A change event [7, Section 13.3.7] is an event which is generated when a certain condition becomes true. The condition typically refers to the states of objects referenced through associations. In Cassandra/xUML, change events are denoted by  $\text{when}(cond)$ , where  $cond$  is a boolean expression. The UML 2.2 semantics [7, Section 13.3.7] for change events is under-specified. For example, it is not indicated when a change event is evaluated or how a change event is detected. Also, implementations may or may not let

change events remain in case their condition becomes false again after having been true.

Given a state machine  $X$  with a transition  $t$  labelled by a change event  $\text{when}(cond)$ , the Cassandra/xUML simulator adds an event  $e_{\text{when}(cond)}$  to the event pool of  $X$  whenever  $cond$  changes from false to true (personal communication with KnowGravity). The event  $e_{\text{when}(cond)}$  triggers transition  $t$  once dispatched and remains in the event pool even in case  $cond$  becomes false before  $e_{\text{when}(cond)}$  is dispatched.

If a dispatched event is not the trigger of an enabled transition, the event is discarded. Otherwise, the actions labelling the transition are carried out. The only type of action we currently allow is the sending of a signal [7, Section 11.3.45], where the target may either be an object within the system or the environment.

#### B. Run-to-completion

An important aspect of concurrency is the interleaving of process executions. Run-to-completion (RTC) assumptions can help reduce the complexity of a concurrent system. A *local RTC step* of a state machine  $X$  consists of processing all actions labelling a transition triggered by some event. In the literature, three different levels of RTC seem to be considered (no fixed terminology seems to exist and the names are our own):

*Local RTC*: A local RTC step of a state machine  $X$  must be completed before the next event can be dispatched to  $X$ .

*Atomic RTC*: While a state machine  $X$  is executing a local RTC step, no other event can be dispatched to any state machine in the system.

*Global RTC*: External signals may only be dispatched to the system in case all event pools are empty and there are no remaining change events.

Local RTC is required by the UML specification [7, Section 15.3.12]. It ensures that a state machine is in a well-defined configuration before the next event is dispatched.

In addition to local RTC, implementations may enforce stricter notions of RTC. Atomic RTC is employed, e.g., by [8] and [9]. With atomic RTC, local RTC steps in different state machines may not be interleaved (which is not forbidden by local RTC). Global RTC separates internal system interactions (between objects) from interactions with the environment. Note that atomic RTC implies local RTC, but that global RTC implies neither local nor atomic RTC. The Cassandra/xUML simulator uses both atomic RTC and global RTC [6, Section 4.3.5].

## IV. TRANSLATION INTO MCRL2

#### A. The mCRL2 language

The mCRL2 specification language [2] extends the process algebra ACP [3] with abstract data types. Built-in data types include booleans, integers, and lists. New structured data types can be defined using the keyword `struct`. We currently only use enumerated data types, e.g.,

sort `Elt_State` = struct `Ready` | `Not_Ready`. Functions over sorts can be defined by giving equations.

The process specification language of mCRL2 allows for the definition of basic actions with zero or more parameters. For example, `act send, read: Message` defines the actions `send` and `read` which take a parameter of type `Message`. Similarly, a process specification may take parameters, e.g., `proc Element(state: Elt_State)`. Processes can be composed using sequential and parallel composition and non-deterministic choice. Actions can be hidden (turned into the silent action) and blocked (disallowed). Synchronisation is achieved by using a communication operator  $\Gamma_C$ , where elements of  $C$  are of the form  $a_1 \mid a_2 \mid \dots \mid a_n \rightarrow c$ , meaning that the action  $c$  is the result of the multi-party synchronisation of the actions  $a_1, a_2, \dots, a_n$ .

An mCRL2 specification consists of data type definitions, equations over the data types, process specifications and an initial process. The above process specification `procElement(state: Elt_State)` could, e.g., be initialised as `init Element(Ready)`.

### B. The translation

In our translation from xUML to mCRL2, each class becomes a process specification. Each of these process specifications consists of two parallel parts: One part is the translation of the state machine associated with the class, the other part formalises the event pool associated with the state machine as a buffer process. The buffer process essentially implements a queue. An event is placed in the queue by a synchronous communication between the sending process and the buffer process. The sending process can either be another process representing a state machine, the environment or a process monitoring change events (described at the end of this section). Signals are dispatched on a FIFO basis through synchronous communication between the buffer process and the process representing a state machine.

*Class diagrams:* As mentioned in Section III-A, we allow for class generalisations and class associations. In our translation, the first is dealt with by flattening the class hierarchy; each superclass  $Y$  of a class  $X$  occurs only once in this flattening, even in case there are several *is-a* associations between  $X$  and  $Y$  in the class diagram. Now, if  $X$  is a class with superclasses  $Y_1, \dots, Y_n$ , the flattened class  $X'$  arising from  $X$  has all attributes of  $X, Y_1, \dots, Y_n$ , and the state machine of  $X'$  is defined as the concurrent composition of the state machines of  $X, Y_1, \dots, Y_n$ .

Class associations are translated by defining an enumerated type consisting of identifiers (depending on the instantiation of the model) and supplementing each mCRL2 representation of a class instance with one parameter (of the enumerated type) for each of its associations.<sup>6</sup> For example,

<sup>6</sup>In practice we employ macro pre-processing of the mCRL2 specification before model checking. This is to avoid loop constructs when dealing with one-many and many-many associations.

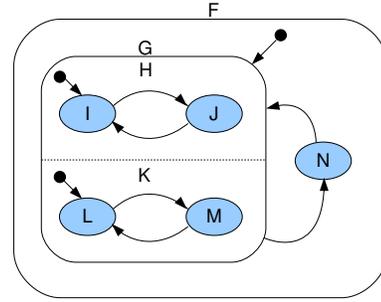


Figure 1. A state machine

if each instance of a class  $X$  is associated with exactly one instance of a class  $Y$ , then the process specification of  $X$  will be of the form `proc X(..., id_Y: Id, ...)`, where `Id` is the enumerated type consisting of identifiers.

*State machines:* The potential state configurations of a state machine  $X$  are encoded as follows. For each non-concurrent composite state (OR-state)  $S$ , we define an enumerated type `ancS_states` where `ancS` identifies  $S$  in the state hierarchy. If  $S$  has substates  $P, \dots, Q$ , then `ancS_states` has members `ancS_substate_P, \dots, ancS_substate_Q`, and `ancS_substate_nop`. The process specification of the state machine  $X$  is then supplied with a parameter `ancS_state` whose value represents the currently active substate of  $S$ . In particular, the top state  $T$  of a state machine for a class is always a non-concurrent state and gives rise to a parameter `T_state`. We refer to `ancS_state` as a *state parameter*. If  $S$  is not active, then `ancS_state` has the value `ancS_substate_nop`.

The configurations of a concurrent state  $S$  are not modelled by parameters, as they are determined by the state configurations of the (direct) substates of the concurrent components of  $S$  (the Cartesian product of these substates to be precise). To illustrate, consider the state machine  $F$  in Figure 1 (transitions are unlabelled as we only wish to illustrate how composite states are treated). In Figure 2 we list the data type definitions arising from  $F$  together with the declaration of state parameters in the process specification of  $F$  (disregarding any class attributes), and an initial process corresponding with the initial state configuration.

Since we treat class generalisation by flattening, if a class  $Y$  generalises a class  $X$ , then the process specification of the state machine for  $X$  (which concurrently composes the state machines for  $Y$  and  $X$ ) will have the state parameters arising from both  $X$  and  $Y$ . This is completely analogous to the handling of concurrent states within a state machine.

*Transitions:* Following the local RTC assumption, a process specifying the state machine of a class  $X$  can obtain an event from its buffer process (event pool) precisely when it is in a stable state, i.e., when no other transition is currently being taken. One of the transitions triggered by

```

sort F_states      = struct F_substate_G
                    | F_substate_N
                    | F_substate_nop;
sort F_G_H_states = struct F_G_H_substate_I
                    | F_G_H_substate_J
                    | F_G_H_substate_nop;
sort F_G_K_states = struct F_G_K_substate_L
                    | F_G_K_substate_M
                    | F_G_K_substate_nop;

proc F(F_state: F_states,
      F_G_H_state: F_G_H_states,
      F_G_K_state: F_G_K_states) = ...;

init F(F_substate_G,
      F_G_H_substate_I,
      F_G_K_substate_L);

```

Figure 2. Translation of the state machine in Figure 1: data types for representing states, state parameters and initialisation

the obtained event is taken at random (non-deterministic choice), assuming such a transition exists. The actions labelling the chosen transition are executed and the state parameters of the process are updated to reflect the new state configuration. If no transition is triggered by the event, then the event is discarded, as allowed by the UML state machine semantics [7, Section 15.3.12].

*Change events:* We implement change events by introducing a process for each such event. This process monitors the value of the condition in the change event. For example, if a state machine  $X$  has a transition from a state  $S$  triggered by (in pseudo-notation)  $\text{when}(P.\text{state} = T \ \& \ Q.\text{state} = U)$ , where  $P$  and  $Q$  are state machines associated with  $X$ , then we create a process  $\text{when\_X\_S}(P\_in\_state\_T : \text{Bool}, Q\_in\_state\_U : \text{Bool})$ , and let the objects  $P$  and  $Q$  communicate synchronously with the monitor process whenever they enter or leave the states  $T$  and  $U$ , respectively. This communication updates the values of the boolean parameters  $P\_in\_state\_T$  and  $Q\_in\_state\_U$ . When an update results in the condition changing from false to true, the monitor process places a signal in the buffer of  $X$ . This signal remains in the buffer of  $X$  even when the condition becomes false again. Hence, at the moment the state machine reacts to the event, the condition might no longer hold.

*Architecture:* We summarise how the elements of an xUML model are mapped onto the elements of an mCRL2 specification: For each flattened class  $X$  with associated state machine  $S$ , we define a process specification  $\text{proc } X(\text{id} : \text{Id}, \dots)$  consisting of the parallel composition of  $S(\text{id} : \text{Id}, \dots)$ ,  $S\_buffer(\text{id} : \text{Id}, \dots)$ , and  $S\_when_i(\text{id} : \text{Id}, \dots)$  where  $i$  ranges over the change events of  $S$ . The parameter  $\text{id}$  represents the unique identifier associated with an instance of the represented class. The translation of the state machine  $S$  is embodied by  $S$  and  $S\_buffer$  represents the event pool. Each  $S\_when_i$  monitors one of the change events occurring in the state machine

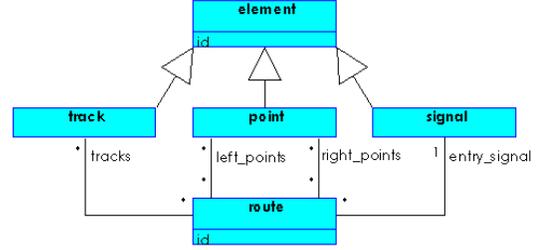


Figure 3. Class diagram for Micro interlocking

associated with  $X$ , as described earlier.

An instance of an xUML model defines, in addition to the above, an enumerated type with object identifiers and an appropriate initialisation consisting of the parallel composition of the processes arising from the objects in the instantiation, together with the required synchronisation constraints.

## V. PROBLEMS AND POSSIBLE SOLUTIONS

The above translation presents us with the following problems from a model checking perspective:

- 1) The system has an infinite state space as there is no bound on the number of messages accepted by the system from the environment.
- 2) System components may starve, i.e., a process may have events in its buffer waiting to be dispatched, but the process may never get its turn.

To alleviate these problems, we propose two constraints:

- A) Limit the size of buffers. This restriction will overcome the first of the above problems. However, as an object may send several signals to itself during a local RTC step (cf. Section III-B), we only impose this limit on messages coming from other objects and from the environment in order to avoid deadlock.
- B) Add a mechanism which ensures that the system can only receive a message from the environment in case all message queues are empty. In other words, implement global RTC (cf. Section III-B). This restriction addresses both 1 and 2 under the assumption that external signals do not (directly or indirectly) trigger an unbounded number of internal events. Consequently, each process will eventually get the chance to empty its buffer.

## VI. MODEL CHECKING A TOY SPECIFICATION

We have applied our translation to a toy interlocking specification which we refer to as the Micro model. The Micro model has classes named *element*, *track*, *point*, *signal* and *route*, where *element* is a generalisation of *track*, *point* and *signal*. The class diagram for this model is depicted in Figure 3. An instance of the Micro model is obtained from the track layout depicted in Figure 4.

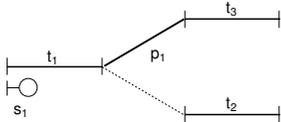


Figure 4. Track layout for an instance of the Micro model

The layout consists of three tracks  $t_1$ ,  $t_2$ ,  $t_3$ , one point  $p_1$  (positioned left in the picture), one signal  $s_1$ , and two routes: route  $r_1$  requires  $p_1$  to be positioned left and goes from track  $t_1$  to track  $t_3$ ; route  $r_2$  requires  $p_1$  to be positioned right and goes from  $t_1$  to  $t_2$ ; both routes have  $s_1$  as their entry signal. The model instance thus contains three track objects, one point object, one signal object and two route objects, and these objects are linked via the following associations:

$$\begin{aligned}
 \text{tracks} &= \{ \langle r_1, t_1 \rangle, \langle r_1, t_3 \rangle, \langle r_2, t_1 \rangle, \langle r_2, t_2 \rangle \} \\
 \text{left\_points} &= \{ \langle r_1, p_1 \rangle \} \\
 \text{right\_points} &= \{ \langle r_2, p_1 \rangle \} \\
 \text{entry\_signal} &= \{ \langle r_1, s_1 \rangle, \langle r_2, s_1 \rangle \}
 \end{aligned}$$

The main functionality of the Micro model is route setting and route cancellation. Informally described, when a route receives a reserve request, it should signal to its left points and right points to move into position. When all points are positioned, all tracks along the route are clear and the entry signal is ready, the entry signal is set to show proceed. When one of the elements associated with the route is no longer in the required state, or the route is cancelled, the route entry signal is set to show stop. The state machine describing the behaviour of the class Route is shown in Figure 5.

We translated the Micro model instance into mCRL2 in two versions corresponding to constraints A and B from Section V. The state space resulting from version A is huge even with buffer size 1 ( $61 \times 10^{12}$  states), but our symbolic tools still compute the number of states within seconds: The mentioned state space was explored in 113 seconds using 238 MB memory of a machine equipped with an Intel Xeon 2.66 GHz, 32GB of memory and Linux 2.6.18. To obtain version B we used barrier synchronisation. This version has a significantly smaller state space (8 million states). However, computing the number of states takes longer (160 seconds). The state space reduction that stems from a global RTC assumption is thus significant, but run-time increases.

We were able to prove the presence of certain (unwanted) traces in both version A and B by placing a monitor process in parallel with the mCRL2 translation. The monitor deadlocks the process in case a certain finite trace, representing the violation of a safety property, is present. The deadlock detection functionality of the symbolic model checker from the LTSmin toolset was used to produce a trace. The trace shows that when the entry signal  $s_1$  has been set to show proceed, the system may command the point  $p_1$  to move

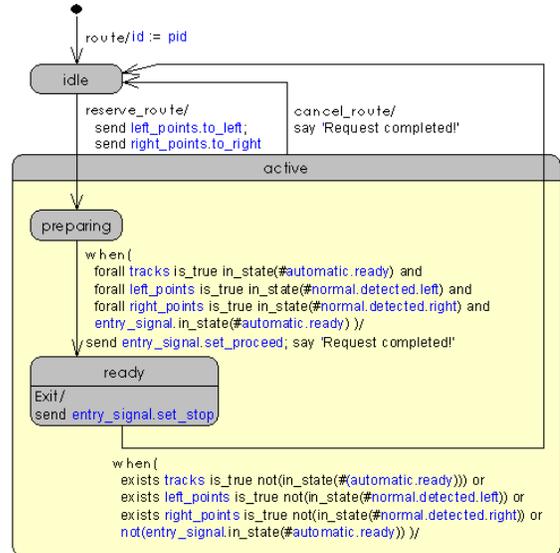


Figure 5. State machine for class Route

before it sets  $s_1$  to show stop, thus risking the derailment of a train passing over  $p_1$ . However, we point out that the Micro model is merely intended to illustrate the type of xUML model constructs that are used in the modelling of interlockings, and it is not claimed to be a safe interlocking specification. Our main point is that some of these bad traces cannot be observed in the Cassandra/xUML simulator, because it uses a stronger RTC assumption than our versions A (only local RTC) and B (local and global RTC).

## VII. RELATED WORK

There is extensive work on the formalisation of executable UML, and in particular, UML state machines, for the purpose of carrying out formal verification. See for example, [8–12]. Translation of xUML into a process algebraic language has been done in [13], [14]. More references can be found in [11] and the survey article [15]. In the above work, translation focuses on composite and concurrent states, history pseudo-states, (conflicting) transitions, and the action language. We were not able to locate any research that includes the formalisation of change events, which are an essential ingredient in the specification of high-level interlocking requirements.

Formal methods have been widely applied in the verification of interlockings. The work can be divided into two categories: Verification of concrete interlockings (e.g. [16–18]) and verification of more generic descriptions (e.g. [19], [20]), as in our case.

## VIII. CONCLUSION

We have presented a translation of a subset of xUML into the process algebra mCRL2. Each of the elements of the xUML subset is translated into a very simple mCRL2

construct, with the translation of the change events being the most complicated. We expect it to be straightforward to extend our translation to include transitions with guards, as well as history and final pseudo-states. We also expect the constructs of any chosen action language to be included easily, bar operations like object creation and destruction. These last operations would correspond to on-the-fly process creation and destruction which is not possible in mCRL2. Given the simplicity of the mCRL2 constructs, we expect that the presented translation can be automated without too much trouble. In fact, work on this automatic translation from xUML (in the form of XMI files) to mCRL2 has already begun.

Our first steps towards verifying safety properties of xUML interlocking specifications have demonstrated the following: First, the fairness and safety properties of an interlocking system may depend crucially on the run-to-completion (RTC) assumption employed in the implementation. Verification should therefore be relative to a particular choice of RTC.

Second, even for small xUML models, such as our toy specification, the state space can be enormous. Still the symbolic model checker seems to be able to deal quite well with the mCRL2 translations obtained from the toy specification. However, in order to verify real interlocking specifications, we expect it will be necessary to come up with specific abstraction and decomposition techniques, as well as reduce the number of event orderings, either in a generic way (partial-order reduction) or specifically (such as RTC assumptions).

#### REFERENCES

- [1] S. J. Mellor and M. Balcer, *Executable UML: A foundation for model-driven architecture*. Addison Wesley, 2002.
- [2] J. F. Groote, A. Mathijssen, M. A. Reniers, Y. S. Usenko, and M. van Weerdenburg, "The formal specification language mCRL2," in *Proc. of Methods for Modelling Software Systems*, ser. Dagstuhl Seminar Proceedings, vol. 06351, 2007.
- [3] J. A. Bergstra and J. W. Klop, "Process algebra for synchronous communication," *Information and Control*, vol. 60, no. 1-3, pp. 109–137, 1984.
- [4] S. Blom and J. van de Pol, "Symbolic reachability for process algebras with recursive data types," in *Proc. Theoretical Aspects of Computing (ICTAC 2008)*, ser. Lecture Notes in Computer Science, vol. 5160. Springer, 2008, pp. 81–95.
- [5] S. C. C. Blom, J. C. van de Pol, and M. Weber, "Bridging the gap between enumerative and symbolic model checkers," CTIT, University of Twente, Enschede, Technical Report TR-CTIT-09-30, 2009.
- [6] *Cassandra/xUML User's Guide*, KnowGravity Inc., 2008. [Online]. Available: <http://www.knowgravity.com/eng/value/cassandra.htm>
- [7] (2009, Feb.) OMG Unified Modeling Language Superstructure Version 2.2. Object Management Group. [Online]. Available: <http://www.omg.org/spec/UML/2.2/Superstructure>
- [8] M. von der Beeck, "Formalization of UML-statecharts," in *Proc. UML 2001*, ser. Lecture Notes in Computer Science, vol. 2185. Springer, 2001, pp. 406–421.
- [9] Fei Xie, V. Levin, and J. Browne, "Model checking for an executable subset of UML," in *16th IEEE International Conference on Automated Software Engineering (ASE 2001)*, 2001, pp. 333–336.
- [10] R. Alur and M. Yannakakis, "Model checking of hierarchical state machines," *ACM Transactions on Programming Languages and Systems*, vol. 23, no. 3, pp. 273–303, 2001.
- [11] W. Damm, B. Josko, A. Pnueli, and A. Votintseva, "A discrete-time UML semantics for concurrency and communication in safety-critical applications," *Science of Computer Programming*, vol. 55, pp. 81–155, 2005.
- [12] Z. Hu and S. M. Shatz, "Explicit modeling of semantics associated with composite states in UML statecharts," *Journal of Automated Software Engineering*, vol. 13, no. 4, pp. 423–467, Oct. 2006.
- [13] E. Turner, H. Treharne, S. Schneider, and N. Evans, "Automatic generation of CSP || B skeletons from xUML models," in *Proc. of Theoretical Aspects of Computing (ICTAC 2008)*, 2008, pp. 364–379.
- [14] W. L. Yeung, K. R. P. H. Leung, J. Wang, and W. Dong, "Improvements towards formalizing UML state diagrams in CSP," in *Proc. of the 12th Asia-Pacific Software Engineering Conference (APSEC 2005)*. IEEE Computer Society, 2005.
- [15] B. Purandar and S. Ramesh. (2004, Jul.) Model checking of statechart models: Survey and research directions. [Online]. Available: <http://arxiv.org/abs/cs.SE/0407038>
- [16] S. Gnesi, D. Latella, G. Lenzini, C. Abbaneo, A. M. Amendola, and P. Marmo, "An automatic SPIN validation of a safety critical railway control system," in *Proc. of the 2000 Int. Conf. on Dependable Systems and Networks*. Washington, DC, USA: IEEE Computer Society, 2000, pp. 119–124.
- [17] A. Cimatti, F. Giunchiglia, G. Mongardi, D. Romano, F. Torielli, and P. Traverso, "Formal verification of a railway interlocking system using model checking," *Formal Aspects of Computing*, vol. 10, no. 4, pp. 361–380, 1998.
- [18] W. Fokkink, "Safety criteria for the vital processor interlocking at Hoorn-Kersenboogerd," in *5th Conference on Computers in Railways (COMPRAIL'96). Volume I: Railway Systems and Management*, 1996.
- [19] K. Winter and N. J. Robinson, "Modelling large railway interlockings and model checking small ones," in *ACSC '03: Proc. of the 26th Australasian comp. sci. conference*. Australian Computer Society, Inc., 2003, pp. 309–316.
- [20] L.-H. Eriksson, "Specifying railway interlocking requirements for practical use," in *Proceedings of the 15th International Conference on Computer Safety, Reliability and Security (SAFECOMP'96)*. Springer, 1996.