

Model checking of UML models via a mapping to communicating extended timed automata^{*}

Iulian Ober, Susanne Graf, Ileana Ober

VERIMAG
2, av. de Vignate
38610 Gières, France
E-mail: {ober,graf,iober}@imag.fr

Abstract. We present a technique and a tool for model-checking operational UML models based on a mapping of object oriented UML models into a framework of communicating extended timed automata - in the IF format - and the use of the existing model-checking and simulation tools for this format.

We take into account most of the structural and behavioral characteristics of classes and their interplay and tackle issues like the combination of operations, state machines, inheritance and polymorphism, with a particular semantic profile for communication and concurrency. The UML dialect considered here, also includes a set of extensions for expressing timing.

Our approach is implemented by a tool importing UML models via an XMI repository, and thus supporting several commercial and non-commercial UML editors. For user friendly interactive simulation, an interface has been built, presenting feedback to the user in terms of the original UML model. Model-checking and model exploration can be done by reusing the existing IF state-of-the-art validation environment.

1 Introduction

The Unified Modeling Language (UML) provides a means for describing operational models of a system at various levels of abstraction, corresponding to different development phases. UML based model driven development has become rapidly a standard in recent years and supported by several case tools for editing UML models and generating code. In the context of real-time and safety critical systems, model based formal validation is essential, but is today very little supported by UML case tools.

This work is part of the OMEGA IST project, whose aim is building a basis for a UML based development environment for real-time and embedded systems, including a set of notations for different aspects with common semantic foundations, tool supported verification methods for large systems, including real-time related aspects [32]. We present a technique and a tool for model-checking operational UML models based on a mapping from object oriented UML models to a system of communicating extended timed automata, as defined by the IF format and on the use of existing model-checking and simulation tools for this format [7].

^{*} This work is supported by the OMEGA European Project (IST-33522). See also <http://www-omega.imag.fr>

The definition of a verification framework for UML models involves a number of design choices. They concern notably the set of *covered UML language elements*, possible *language extensions* (such as a formalism for specifying properties to be verified), and the *operational semantics* of a model.

In terms of language coverage, we focus on the operational part of UML: classes with structural and behavioral features, relationships (associations, inheritance), behavior descriptions through state machines and actions. The issues we tackle, like the combination of *operations* and *state machines*, *inheritance* and *polymorphism*, *run-to-completion* and *concurrency*, go beyond the previous work done in this area (see section 1.1), which has mainly focused on verification of statecharts. Our choices are outlined in section 2.

The formalisation of the operational semantics of UML models is based on a mapping from UML into an intermediate formal representation IF[6] based on *communicating extended timed automata* (CETA). This choice is motivated by the existence a verification toolset based on this semantic model [7, 9] which has been productively used in a number of research projects and case studies, e.g. in [8, 16]. The main features of the CETA model are presented in section 1.2, and in section 3 we discuss the mapping from UML into this model.

An important issue in designing real-time systems is the ability to capture quantitative timing requirements and assumptions, as well as time dependent behavior. We rely on the extensions defined in the context of the Omega project [17, 15]. We summarize these extensions and their mapping into IF in section 4.

Another important issue is the formalism in which *properties* are expressed, which may range from temporal logics to automata-based specifications. In section 5 we introduce a simple property description language that reuses some concepts from UML (like objects, state machines) while remaining sufficiently expressive for a large class of properties (equivalent to the linear temporal logic LTL). The language of *observer objects* makes use of concepts that are familiar to any UML user, and has the potential to alleviate the cultural shock of introducing formal dynamic verification to UML models.

Finally, section 6 presents the *UML validation toolset*. By using the IF tools as underlying simulation and verification engine, the UML tools presented here benefit from a large spectrum of model reduction and analysis techniques already implemented therein, such as *static analysis* and optimizations for state-space reduction, *partial order* reductions, some forms of *symbolic* exploration, model minimization and comparison, etc [7, 9]. These techniques improve the scalability of model-checking, which is essential when analyzing UML models.

The techniques and the tool presented in this paper are subject to experimental validation on several larger case studies within the OMEGA project [1].

1.1 Related work

The application of formal analysis techniques (and particularly model checking) to UML is a very active field of study in recent years, as witnessed by the number of papers on this subject ([26, 27, 25, 24, 23, 31, 13, 14, 34, 4] are most often cited).

Like ourselves, most of these authors base their work on an existing model checker (SPIN[20] in the case of [26, 27, 25, 31], COSPAN[19] in the case of [34],

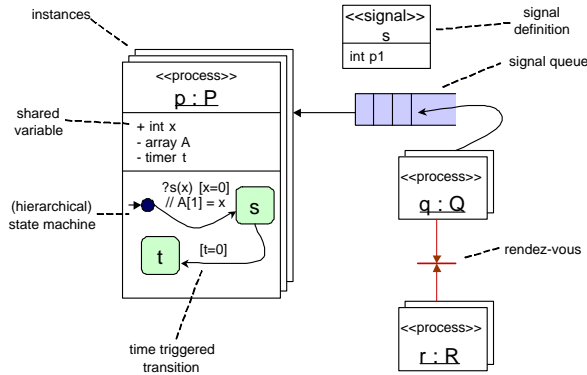


Fig. 1. Constituents of a CETA model.

Kronos[35] for [4] and UPPAAL[22] for [23]), and on the mapping of UML to the input language of the respective tool.

For specifying properties, some authors opt for the property language of the model checker itself (e.g. [25, 26, 27]). Others use UML collaboration diagrams (e.g. [23, 31]) which are too weak to express all relevant properties. We propose to use a variant of UML state machines to express properties in terms of observers.

Concerning language coverage, all previous approaches are restricted to *flat class structures* (no inheritance) and to behaviors, specified *exclusively by statecharts*. In this respect, many important features which make UML an object-oriented formalism (inheritance, polymorphism and dynamic binding of operations) are missed. Our approach is, to our knowledge, the first to try to fill this gap.

Our starting point for handling of UML state machines (not described in detail in this paper) was the material cited above together with previous work on Statecharts ([18, 12, 28] to mention only a few). In the definition of our concurrency model we have taken inspiration from our previous assessment of the UML concurrency model [29], and from other positions on this topic (see for example [33]) and we respected the operational semantics defined in the OMEGA project [11].

1.2 The back-end model, techniques and tools

The validation approach proposed in this work is based on the formal model of communicating extended timed automata (CETA) and on the IF environment built around this model [7, 9, 10]. We summarize the elements of this model in the following.

Modeling with communicating extended automata

The CETA model and IF were developed at VERIMAG in order to provide an instrument for modeling and validating *distributed systems* that can manipulate *complex data*, may involve *dynamic aspects* and *real time constraints*. Additionally, the model allows to describe the semantics of higher level formalisms (e.g. UML or SDL) and has been used as a format for inter-connecting validation tools.

In this model, a system is composed of a set of communicating *processes* that run in parallel (see figure 1). Processes are instances of *process types*. They have their own identity (PID), they may own complex data variables (defined through ADA-like data type definitions), and their behavior is defined by a *state machine*. The state machine of a process type may use composite states and the effect of transitions is described using common (structured) imperative statements.

Processes may inter-communicate via *asynchronous signals*, via *shared variables* or via *rendez-vous*. Parallel processes are composed asynchronously (i.e. by interleaving). The model also allows *dynamic creation* of processes, which is an essential feature for modeling object systems that are by definition dynamic.

The CETA model allows to describe the link between system execution and time progress in a precise manner, and thus offers support for modeling real time constraints. The concepts used by CETA are those of *timed automata with urgency* [3, 5]: there are special variables called *clocks* which measure time progress and which can be used in transition guards. A special attribute of each transition, called *urgency*, specifies how time may progress when the transition is enabled.

A framework for modeling priority

On top of the above model, we use a framework for specifying dynamic priorities via partial orders between processes. The framework was formalized in [2]. Basically, a CETA model is associated with a set of priority directives of the form: (*state condition*) $\Rightarrow p_1 \prec p_2$. They are interpreted as follows: given a system state and a directive, if the condition of the directive holds in that state, then process with ID p_1 has priority over p_2 for the next move (meaning that if p_1 has an enabled transition, then p_2 is not allowed to move).

Property specification with observers

Dynamic properties of CETA/IF models may be expressed using *observer automata*. These are special processes that may monitor¹ the changes in the *state* of a model (variable values, contents of queues, etc.) and the *events* occurring in it (inputs/outputs, creation/destruction of processes, etc.).

For expressing properties, the states of an observer may be classified (syntactically) as *ordinary*, *error* or *success*. Observers may be used to express *safety properties*. A re-interpretation of success states as accepting states of a Büchi automaton could also allow observers to express liveness properties.

CETA/IF observers are rooted in the observer concept introduced by Jard, Groz et Monin in the VEDA tool [21]. This intuitive and powerful property specification formalism has been adapted over the past 15 years to other languages (LOTOS, SDL) and implemented by industrial case tools like Telelogic's ObjectGEODE.

Analysis techniques and the IF-2 toolbox

The IF-2 toolbox [7, 9] is the validation environment built around the model presented before. It is composed of three categories of tools:

1. **behavioral tools** for simulation, verification of properties, automatic test generation, model manipulation (minimization, comparison). The tools implement

¹ The semantics is that observer transitions synchronize with the transitions of the CETA model.

techniques such as *partial order reductions* and *symbolic* simulation of time, and thus present a good level of scalability.

2. **static analysis tools** which provide source-level optimizations that help reducing furthermore the state space of the models, and thus improve the chance of obtaining results from the behavioral tools. Among the state of the art techniques that are implemented we mention *data flow analysis* (e.g. dead variable reduction), *slicing* and simple forms of *abstraction*.
3. **front-ends and exporting tools** which provide source-level coupling to higher-level languages (UML, SDL) and to other verification tools (Spin, Agatha, etc.).

The toolbox has already been used in a series of industrial-size case studies [7, 9].

2 Ingredients of UML models

This section outlines our design choices with respect to *the UML concepts covered* and *the computation* and to the *execution model* adopted.

2.1 UML concepts covered

In this work we consider an operational subset of UML, which includes the following UML concepts: active and passive *classes* - with their *operations* and *attributes*, *associations*, *generalizations* - including polymorphism and dynamic binding of operations, *basic data types*, *signals*, and *state machines*. State machines are not discussed in this paper as they are already tackled in many previous works like [26, 27, 25, 24, 23, 31, 14, 34, 4].

Additionally to the elements mentioned above, a number of UML extensions for describing timing constraints and assumptions are supported. They were introduced in [15, 17] and are discussed in section 4.

2.2 The execution model

We describe in this section some of the design choices made with respect to the computation and the concurrency model implemented by our method and tools. The purpose is to illustrate some of the particularities of the model and not to give a complete/formal semantics for UML. Actually, a precise semantics is given by our proposed mapping to CETA/IF, outlined in the next section and implemented by our tools.

The execution model presented in the following is the one established by the OMEGA project [11]. Nevertheless, other execution models can be accommodated to our framework by adapting the mapping to CETA/IF accordingly.

Activity groups and concurrency. There are two kinds of classes: *active* and *passive*, both being described by attributes, relationships, operations and state machines.

At execution, each instance of an active class defines a concurrency unit called *activity group*. Each instance of a passive class belongs to exactly one activity group.

Different activity groups execute concurrently, and objects inside the same an activity group execute sequentially. This means that requests (signals or operations)

are sequentialized at the border of the activity group, and handled one by one when the whole group is *stable*.

The notion of *stability* is defined as follows: an *object* is stable if it has nothing to execute spontaneously and no pending operation call from inside its group. An *activity group* is stable when all its objects are stable.

The above notion of stability defines a notion of *run-to-completion* step for activity groups: a step is the sequence of actions executed by the objects of the group from the moment an external request is taken by one of the group objects, and until the group becomes stable. During a step, other requests coming from outside the activity group are not handled and are enqueued.

Operations and state machines. We assume two kinds of operations (distinguished syntactically), allowed in any class: *primitive* operations and *triggered* operations. *Primitive operations* have the body described by a method (with an associated action), while *triggered operations* are handled in the state machine of a class, and their effect is described on transitions. Triggered operations differ from *signals* in that they may have a return value.

With respect to concurrency, the two kinds of operations are handled differently. At any moment an object having the control may call a primitive operation on an object from the same activity group, and the call is stacked and handled immediately. On the contrary, triggered operations and signals go to the boundary of the active group and are queued for handling in a later run-to-completion step. Primitive operation calls that transgress the boundary of an active group are also queued and handled like signals and triggered operations.

3 Mapping UML models to CETA/IF

In this section we give the main lines of the mapping of a UML model to a CETA model. The idea is to obtain a CETA model whose operational semantics is identical to the desired semantics of the UML model. This additional layer helps us tackle with the complexity of UML, and provides a semantic basis for re-using our existing model checking tools (see section 6).

The mapping is done in a way that all runtime UML entities (objects, call stacks, pending messages, etc.) are identifiable as a part of the CETA model's state. In simulation and verification, this allows tracing back to the UML specification.

3.1 Mapping the object domain to CETA

Mapping of attributes and associations. Every class X is mapped to a process type P_X that will have a local variable of corresponding type for each attribute or association of X . As inheritance is flattened in the CETA model, all inherited attributes and associations are replicated in the processes corresponding to each class.

Activity group management. Each *activity group* is managed at runtime by a special process of a type called *GM*. This process sequentializes the calls coming from outside the activity group, and helps to ensure the run-to-completion policy. In each P_X there is a local variable *leader*, which points to the *GM* process managing its activity group.

Mapping of operations and call polymorphism. For each operation $m(p_1 : t_1, p_2 : t_2, \dots)$ in class X , the following components are defined in the CETA model:

- a signal $call_{X::m}(waiting : pid, caller : pid, callee : pid, p_1 : t_1, p_2 : t_2, \dots)$ used to indicate an operation call. If the call is made in the same activity group, *waiting* indicates the process that waits for the completion of the call in order to continue execution. *caller* designates the process that is waiting for a return value, while *callee* designates the process corresponding to the object receiving the call (a P_X instance).
- a signal $return_{X::m}(r_1 : tr_1, r_2 : tr_2, \dots)$ used to indicate the return of an operation call (sent to the *caller*). Several return values may be sent with it.
- a signal $complete_{X::m}()$ used to indicate completion of computation in the operation (may differ from return, as an operation is allowed to return a result and continue computation). This signal is sent to the *waiting* process (see $call_{X::m}$).
- if the operation is *primitive* (see 2.2), a process type $P_{X::m}(waiting : pid, caller : pid, callee : pid, p_1 : t_1, p_2 : t_2, \dots)$ which will describe the behavior of the operation using a CETA automaton. The parameters have the same meaning as in the $call_{X::m}$ signal. The *callee* PID is used to access local attributes of the called object, via the shared variable mechanism of CETA.
- if the operation is *triggered* (see 2.2), its implementation will be modeled in the state machine of P_X (see the respective section below). Transitions triggered by a $X :: m$ call event in the UML state machine will be triggered by $call_{X::m}$ in the CETA automaton.

The action of invoking an operation $X :: m$ is modeled in CETA by the sending of a signal $call_{X::m}$. The signal is sent either directly to the concerned object (if the caller is in the same group) or to the object's *active group manager* (if the caller is in a different group). The group manager will enqueue the call and will forward it to the destination when the group becomes stable.

The handling of incoming calls is simply modeled by transition loops (in every state² of the process P_X) which, upon reception of a $call_{X::m}$ will create a new instance of the automaton $P_{X::m}$ and wait for it to finish execution (see sequence diagram in figure 2).

The above mapping provides a simple solution for handling *polymorphic* calls in an inheritance hierarchy: if A and B are a class and its heir, both implementing the method m , then P_A will respond to $call_{A::m}$ by creating a handler process $P_{A::m}$, while P_B will respond to both $call_{A::m}$ and $call_{B::m}$, in each case creating a handler process $P_{B::m}$ (figure 3).

This solution is similar to the one used in most object oriented programming language compilers, where a "method lookup table" is used for dynamic binding of calls to operations; here, the object's state machine plays the role of the lookup table.

Mapping of constructors. Constructors (take $X :: m$ in the following) differ from primitive operations in one respect: their binding is static. As such, they do not need the definition of the $call_{X::m}$ signal and the call (creation) action is directly the creation of the handler process $P_{X::m}$. The handler process begins by creating

² This is eased by the fact that CETA/IF support hierarchical automata.




Fig. 2. Handling primitive operation calls using dynamic creation.




Fig. 3. Mapping of primitive operations and inheritance.

a P_X object and its strong aggregates, after which it continues execution like a normal operation.

Mapping of signals and state machines. UML signals are mapped to signals of the CETA model. UML state machines are mapped almost directly in CETA state machines. Certain transformations are necessary in order to support features that are not directly in the CETA model, such as entry/exit actions, fork/join nodes, history, etc.

Several prior research results tackle the problem of mapping UML statecharts to (hierarchical) automata (e.g. [28]). The method we apply is similar to such approaches.

Actions. The action types supported in the original UML model are *assignments*, *signal output*, *control structure actions*, *object creation*, *method call* and *return*. Some are directly mapped to their CETA/IF counterparts, while the others are mapped as mentioned above to special signal emissions (*call*, *return*) or process creations.

3.2 Modeling run to completion and the activity group concept with dynamic priorities

We discuss here how the concurrency model introduced in section 2.2 is realized in CETA using the dynamic partial priority order mechanism presented in 1.2.

As mentioned, the calls or signals coming from outside an activity group are queued at the border of the group and handled one by one in run-to-completion

steps. In the CETA model, the group management objects (*GM*) handle the simple enqueueing and forwarding behavior.

In order to obtain the desired run-to-completion (RTC), the following priority protocol is applied (the rules concern processes representing instances of UML classes, and not the processes representing operation handlers, etc.):

- All objects of a group have higher priorities than their group manager:
 $\forall x, y. (x.\text{leader} = y) \Rightarrow x \prec y$
 This ensures that as long as an object inside the group may move, the group manager will not initiate the next RTC step.
- Each *GM* object has an attribute *running* which points to the presently or most recently running object in the group. This attribute behaves like a token that is taken or released by the objects having something to execute. The following priority rule:
 $\forall x, y. (x = y.\text{leader.running}) \wedge (x \neq y) \Rightarrow x \prec y$
 ensures that as long as an object has something to execute (the continuation of an action, or the initiation of a new spontaneous transition), no other object in the group may run.
- Every object *x* with the behavior described by a statechart in UML will execute $x.\text{leader.running} := x$ at the beginning of each transition. In regard of the previous rule, such a transition is executed only when the previously running object of the group has reached a stable state, which means that the current object may take the *running* token safely.

4 UML extensions for capturing timing

In order to build a faithful model of a *real-time* system in UML, one needs to represent two types of timing information:

Time-triggered behavior (*prescriptive modeling*): this corresponds, for example, to the common practice in real-time programming environments to link the execution of an action to the expiration of a delay (represented sometimes by a *timer* object).

Knowledge about the timing of events (*descriptive modeling*): information taken as a hypothesis under which the system works. Examples are the worst case execution times of system actions, scheduler latency, etc.

In addition to that, a high-level UML model may also contain timing *requirements* to be imposed upon the system.

Different UML tools targeting real-time systems adopt different UML extensions for expressing such timing information. A standard UML Real-Time Profile, defined by the OMG [30], provides a common set of concepts for modeling timing, but their definition remains mostly syntactic.

We base our work on the framework defined in [17] for modeling timed systems. The framework reuses some of the concepts of the standard real-time profile [30] (e.g. timers, certain data types), and additionally allows expressing *duration constraints* between various events occurring in the system.

4.1 Validation of timed specifications

In this section we present the main concepts taken from [17], that we use in our framework, and we give the principles of their mapping to CETA.

For modeling *time-triggered behavior*, we are using *timer* and *clock* objects compatible with those of [30]. *Clocks* exist natively in the CETA model, and *timers* may be simulated using a clock and a manager process.

The modeling of the *descriptive timing information* makes intensive use of the **events** occurring in a UML system execution. An event has an occurrence time, a type and a set of related information depending on its type. The event types that can be identified are listed in section 5.2, as they also constitute an essential part of our property specification language (presented in section 5). All these UML *events*, and their associated parameters, can be identified in the CETA model. For example: the UML event of invoking an operation $X :: m$ equates to the CETA event of sending the $call_{X::m}$ signal, etc.

If several events of the same type and with the same parameters may occur during a run, there are mechanisms for identifying the particular event occurrence that is relevant in a certain context.

Between the events identified as above, we may define **duration constraints**. The constraints may be either *assumptions* (hypotheses to be enforced upon the system runs) or *assertions* (properties to be tested on system runs).

The class diagram example in figure 4 shows how these events and duration constraints may be used in a UML model. This model describes a typical client-server architecture in which worker objects on the server are supposed to expire after a fixed delay of 10 seconds. A timing assumption attached to the client says that: *"whenever a client connects to the server, it will make a request before its worker object expires, that is before 10 seconds"*.

In a CETA context, for testing or enforcing a timing constraint from the UML model, we are presented with two alternatives. The first alternative corresponds to the case when the constraint is *local* to a CETA process, in the sense that all involved events are directly observed by the process. (For example, the outputs and inputs of a process are directly observed by itself, but they are not visible to other processes.) This is the case in figure 4. In this case, the constraint may be tested or enforced by the CETA process itself, using an additional clock for measuring the duration concerned by the constraint, as well as a CETA transition with an appropriate guard on that clock.

In the second alternative the constraint is not local to a CETA process (we call it *global*). In that case, the constraint will be tested or enforced by a CETA observer running in parallel with the system.

The tools will ensure that runs not satisfying a constraint are either ignored – if it is an assumption, or diagnosed as error – if it is an assertion.

5 Dynamic properties written as UML observers

We discuss in this section a technique for specifying and verifying dynamic properties of UML models, that we call *UML observers*. Similarly to CETA observer automata (section 1.2), UML observers are special objects which run in parallel with a UML system and monitor its *state* and the *events* that occur.

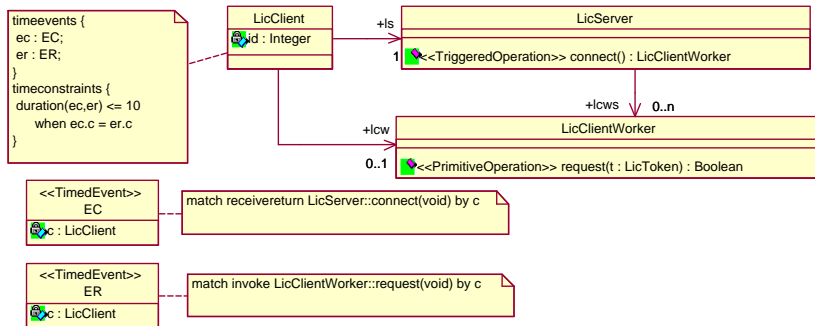


Fig. 4. Using events to describe timing constraints.




Fig. 5. Example of observer for a safety property.

Syntactically, observers are described by special UML classes stereotyped with `<<observer>>`. They may own attributes and methods, and may be created dynamically. An important part of the observer is its *state machine*, which is triggered by events occurring in the UML model, as we will see in the following. The main issue in defining UML observers is the choice of visible event types (which include specific UML event types like operation invocation, etc.).

For UML users, the advantage of UML observers compared to other property specification languages is that they use concepts that are known to UML designers (event driven state machines) while remaining sufficiently formal and expressive.

5.1 An example of property

Let us take a simple example: assume that we have a point-to-point communication protocol described in UML. Two interfaces *TX* and *RX* encapsulate the transmission and reception operations, and, to simplify, at runtime there exists exactly one object implementing each interface. The interface *TX* has one blocking operation *put(p : Data)* (where *Data* is the packet type) and the interface *RX* has one blocking operation *get()* that returns a *Data*.

Assume that we want to express the following reliability property: *whenever put is called with some Data, within at most 5 time units the same Data is received at the other end*. This also supposes that the user at the other end has called *get* within this time frame, reception being signified by the return from *get*. This property is specified in the observer in figure 5.

5.2 Basic observer ingredients

An important ingredient of the observer in figure 5 are the event specifications on some transitions. Here, the notion of event and the event types are the ones introduced in [17]:

- Events related to *operation calls*: **invoke**, **receive** (reception of call), **accept** (start of actual processing of call – may be different from **receive**), **invoke-return** (sending of a return value), **receivereturn** (reception of the return value), **acceptreturn** (actual consumption of the return value).
- Events related to *signal exchange*: **send**, **receive**, **consume**.
- Events related to *actions or transitions*: **start**, **end** (of execution).
- Events related to *states*: **entry**, **exit**.
- Events related to *timers* (this notion is specific to the model considered in [15, 17] and in this work): **set**, **reset**, **occur**, **consume**.

The trigger of an observer transition may be a **match** clause, in which case the transition will be triggered by certain types of events occurring in the UML model. The clause specifies the type of event (e.g. **receive** in figure 5), some related information (e.g. the operation name $TX :: put$) and observer variables that may receive related information (e.g. m which receives the value of the *Data* parameter of *put* in the concerned call).

Besides events, an observer may access any part of the state of the UML model: object attributes and state, signal queues.

As in CETA/IF observers, properties are expressed by classifying observer states as **error**, **success** or **ordinary**.

Writing timing properties. Certain timing properties may be expressed directly in a UML model using the extensions presented in section 4. However, more complicated properties which involve several events and more arbitrary ordering between them may be written using observers. In order to express quantitative timing properties, observers may use the concepts available in our extension of UML, such as *clocks*.

6 The simulation and verification toolset

The principles presented in the previous sections are being implemented in the UML-IF validation toolbox³, the architecture of which is shown in figure 6. With this tool, a designer may simulate and verify UML models and observers developed in third-party editors⁴ and stored in XMI⁵ format. The functionality offered by the

³ See <http://www-verimag.imag.fr/ober/IFx>.

⁴ Rational Rose, I-Logix Rhapsody and Argo UML have been tested for the moment.

⁵ XMI 1.0 or 1.1 for UML 1.4

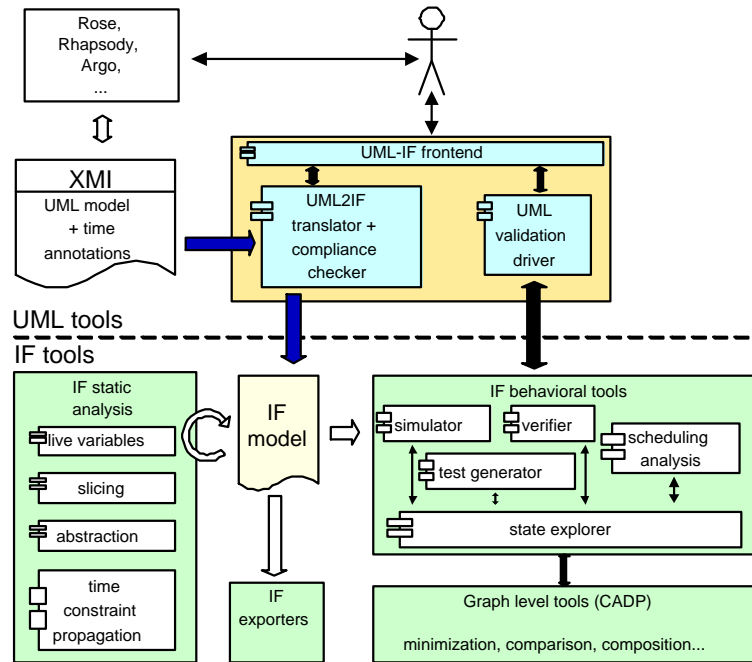


Fig. 6. Architecture of the UML-IF validation toolbox.

tool, is that of an advanced debugger (with step-back, scenario generation, etc.) doubled by a model checker for properties expressed as observers.

In a first phase, the tool generates an IF specification and a set of IF observers corresponding to the model. In a second phase, it drives the IF simulation and verification tools so that the validation results fed back to the user may be marshaled back to level of the original model. Ultimately, the IF back-end tools shall be invisible to the UML designer.

As mentioned in the introduction, by using the IF tools as underlying engine, the UML tools have access to several model reduction and analysis techniques already implemented. Such techniques aim at improving the scalability of the tools, essential in a UML context. Among them, it is worth mentioning *static analysis* and optimizations for state-space reduction, *partial order* reductions, some forms of *symbolic* exploration, model minimization and comparison [7, 9].

A first version of this toolset exists and is currently being used on several case studies in the context of the OMEGA project.

7 Conclusions and plans for future work

We have presented a method and a tool for validating UML models by simulation and model checking, based on a mapping to an automata-based model (communicating extended timed automata).

Although this problem has been previously studied [13, 26, 25, 24, 23, 31], our approach introduces a new dimension by considering the important object-oriented features present in UML: inheritance, polymorphism and dynamic binding of operations, and their interplay with statecharts. We give a solution for modeling these concepts with automata: operations are modeled by dynamically created automata, and thus call stacks are implicitly represented by chains of communicating automata. Dynamic binding is achieved through the use of signals for operation invocation. We also give a solution for modeling run-to-completion and a chosen concurrency semantics using dynamic priorities.

Our experiments on small case studies show that the simulation and model checking overhead introduced by modeling these object-oriented aspects remains low, thus not hampering the scalability of the approach.

For writing and verifying dynamic properties, we propose a formalism that remains within the framework of UML: observer objects. We believe this is an important issue for the adoption of formal techniques by the UML community. Observers are a natural way of writing a large class of properties (linear properties with quantitative time).

The plans for future work include the following main directions:

- assessment of the applicability of our technique to larger models: the tool is beginning to be applied to a set of four case studies provided by industrial partners in the OMEGA project.
- extension of the language scope covered by the tool: we plan to integrate the component and architecture specification framework defined in OMEGA.
- improvement of the ergonomics and integration of the toolset (e.g. the presentation of validation results in terms of the UML model).

References

- [1] <http://www-omega.imag.fr> - website of the IST OMEGA project.
- [2] K. Altisen, G. Gössler, and J. Sifakis. A methodology for the construction of scheduled systems. In M. Joseph, editor, *proc. FTRTFT 2000*, volume 1926 of *LNCS*, pages 106–120. Springer-Verlag, 2000.
- [3] R. Alur and D.L. Dill. A theory of timed automata. In *TCS94*, 1994.
- [4] Vieri Del Bianco, Luigi Lavazza, and Marco Mauri. Model checking UML specifications of real time software. In *Proceedings of 8th International Conference on Engineering of Complex Computer Systems*. IEEE, 2002.
- [5] S. Bornot and J. Sifakis. An algebraic framework for urgency. *Information and Computation*, 163, 2000.
- [6] M. Bozga, J.Cl. Fernandez, L. Ghirvu, S. Graf, J.P. Krimm, and L. Mounier. IF: An intermediate representation and validation environment for timed asynchronous systems. In *Proceedings of Symposium on Formal Methods 99, Toulouse*, number 1708 in *LNCS*. Springer Verlag, September 1999.
- [7] M. Bozga, S. Graf, and L. Mounier. IF-2.0: A validation environment for component-based real-time systems. In *Proceedings of Conference on Computer Aided Verification, CAV'02, Copenhagen*, *LNCS*. Springer Verlag, June 2002.
- [8] M. Bozga, D. Lesens, and L. Mounier. Model-Checking Ariane-5 Flight Program. In *Proceedings of FMICS'01 (Paris, France)*, pages 211–227. INRIA, 2001.

- [9] Marius Bozga, Susanne Graf, and Laurent Mounier. Automated validation of distributed software using the IF environment. In *2001 IEEE International Symposium on Network Computing and Applications (NCA 2001)*. IEEE, October 2001.
- [10] Marius Bozga and Yassine Lakhnech. IF-2.0 common language operational semantics. Technical report, 2002. Deliverable of the IST Advance project, available from the authors.
- [11] W. Damm, B. Josko, A. Pnueli, and A. Votintseva. Understanding UML: A formal semantics of concurrency and communication in real-time UML. In *Proceedings of FMCO'02*, LNCS. Springer Verlag, November 2002.
- [12] Werner Damm, Bernhard Josko, Hardi Hungar, and Amir Pnueli. A compositional real-time semantics of STATEMATE designs. *Lecture Notes in Computer Science*, 1536:186–238, 1998.
- [13] Alexandre David, M. Oliver Möller, and Wang Yi. Formal Verification of UML Statecharts with Real-Time Extensions. In R.-D. Kutsche and H. Weber, editors, *Fundamental Approaches to Software Engineering (FASE'2002)*, volume 2306 of LNCS, pages 218–232. Springer-Verlag, April 2002.
- [14] Maria del Mar Gallardo, Pedro Merino, and Ernesto Pimentel. Debugging UML designs with model checking. *Journal of Object Technology*, 1(2):101–117, August 2002. ([http://www.jot.fm/issues/issue 2002 07/article1](http://www.jot.fm/issues/issue%2002%2007/article1)).
- [15] S. Graf and I. Ober. A real-time profile for UML and how to adapt it to SDL. In *Proceedings of SDL Forum 2003 (to appear)*, LNCS, 2003.
- [16] Susanne Graf and Guoping Jia. Verification experiments on the MASCARA protocol. In *Proceedings of SPIN Workshop '01 (Toronto, Canada)*, January 2001.
- [17] Susanne Graf, Ileana Ober, and Iulian Ober. Timed annotations with UML. Submitted to SVERTS'2003, 2003.
- [18] David Harel and Amnon Naamad. The STATEMATE semantics of statecharts. *ACM Transactions on Software Engineering and Methodology*, 5(4):293–333, 1996.
- [19] Z. Har'El and R. P. Kurshan. Software for Analysis of Coordination. In *Conference on System Science Engineering*. Pergamon Press, 1988.
- [20] G. J. Holzmann. The model-checker SPIN. *IEEE Trans. on Software Engineering*, 23(5), 1999.
- [21] C. Jard, R. Groz, and J.F. Monin. Development of VEDA, a prototyping tool for distributed algorithms. *IEEE Transactions on Software Engineering*, 14(3):339–352, March 1988.
- [22] H. Jensen, K.G. Larsen, and A. Skou. Scaling up UPPAAL: Automatic verification of real-time systems using compositionality and abstraction. In *FTRTFT 2000*, 2000.
- [23] Alexander Knapp, Stephan Merz, and Christopher Rauh. Model checking timed UML state machines and collaborations. In W. Damm and E.-R. Olderog, editors, *7th Intl. Symp. Formal Techniques in Real-Time and Fault Tolerant Systems (FTRTFT 2002)*, volume 2469 of *Lecture Notes in Computer Science*, pages 395–414, Oldenburg, Germany, September 2002. Springer-Verlag.
- [24] Gihwon Kwon. Rewrite rules and operational semantics for model checking UML statecharts. In Bran Selic Andy Evans, Stuart Kent, editor, *Proceedings of UML'2000*, volume 1939 of *Lecture Notes in Computer Science*. Springer-Verlag, 2000.
- [25] D. Latella, I. Majzik, and M. Massink. Automatic verification of a behavioral subset of UML statechart diagrams using the SPiN model-checker. *Formal Aspects of Computing*, (11), 1999.
- [26] J. Lilius and I.P. Paltor. Formalizing UML state machines for model checking. In Rumpe France, editor, *Proceedings of UML'1999*, volume 1723 of *Lecture Notes in Computer Science*. Springer-Verlag, 1999.
- [27] Johan Lilius and Ivan Porres Paltor. vUML: A tool for verifying UML models. In *Proceedings of 14th IEEE International Conference on Automated Software Engineering*. IEEE, 1999.

- [28] Erich Mikk, Yassine Lakhnech, and Michael Siegel. Hierarchical automata as a model for statecharts. In *Proceedings of Asian Computer Science Conference*, volume 1345 of *LNCS*. Springer Verlag, 1997.
- [29] Iulian Ober and Ileana Stan. On the concurrent object model of UML. In *Proceedings of EUROPAR'99*, LNCS. Springer Verlag, 1999.
- [30] OMG. Response to the OMG RFP for Schedulability, Performance and Time, v. 2.0. OMG document ad/2002-03-04, March 2002.
- [31] Timm Schäfer, Alexander Knapp, and Stephan Merz. Model checking UML state machines and collaborations. *Electronic Notes in Theoretical Computer Science*, 55(3):13 pages, 2001.
- [32] writers: Susanne Graf The Omega consortium and Jozef Hooman. The omega vision and workplan. Technical report, Omega Project deliverable, 2003.
- [33] WOODDES. Workshop on concurrency issues in UML. Satelite workshop of UML'2001. See <http://wooddes.intranet.gr/uml2001/Home.htm>.
- [34] Fei Xie, Vladimir Levin, and James C. Browne. Model checking for an executable subset of UML. In *Proceedings of 16th IEEE International Conference on Automated Software Engineering (ASE'01)*. IEEE, 2001.
- [35] S. Yovine. KRONOS: A verification tool for real-time systems. *Springer International Journal of Software Tools for Technology Transfer*, 1(1-2), December 1997.