

REQUIREMENTS-LEVEL SEMANTICS FOR UML STATECHARTS*

Rik Eshuis[†]

Roel Wieringa

*University of Twente, Dept. of Computer Science,
P.O. Box 217, 7500 AE, Enschede, The Netherlands*
{eshuis,roelw}@cs.utwente.nl

Abstract We propose a formal real-time semantics for UML statecharts aimed at the requirements level. A requirements-level model assumes perfect technology and has a considerably simpler semantics than an implementation level model. Our semantics is an adaptation of the STATEMATE statechart semantics, with local variables, real time, identifier addressing, point-to-point communication, synchronous communication and dynamic object creation and deletion. We start with an informal comparison of STATEMATE and UML statechart semantics and then give a formalisation of our semantics in terms of labelled transition systems.

Keywords: statecharts, UML, formal semantics

1. INTRODUCTION

Statecharts were introduced by Harel [5] to model the behaviour of activities in the structured analysis approach STATEMATE [8]. They have been adapted in many object-oriented design notations, including the UML [14], but with an informally or undefined semantics that appears to be quite different from the STATEMATE semantics.

In this paper, we define a formal semantics for UML statecharts based upon the STATEMATE semantics for two reasons. First, a formal semantics allows model checking UML models used for requirements specification. This differs from similar research done for STATEMATE [3] in that it works with OO models and it differs from similar work [11] with UML in that it is at the requirements level, not at the implementation level.

*Partially supported by Esprit Working Group ASPIRE, contract number 22704.

[†]Supported by NWO/SION, grant nr. 612-62-02 (DAEMON).

Table 1 Differences between STATEMATE and requirements-level UML

<i>STATEMATE activity chart + statechart</i>	<i>UML class diagram + statechart</i>
separation of data state/process and control state/process	encapsulation of data state/process and control state/process
channel addressing	identifier addressing
broadcast	point-to-point
instance level model	type-instance distinction

(A requirements-level model focusses on the design and assumes perfect technology [13], whereas an implementation-level model does not, since it is especially concerned with implementing a design.) Second, such a semantics allows to classify the differences between structured and OO semantics of statecharts in two dimensions: structured versus OO models and requirements level versus implementation level.

Tables 1 summarises the difference between our OO semantics and the structured STATEMATE semantics. Firstly, object-oriented models encapsulate data manipulation, control and data state into objects (as operations, statecharts, attributes). Structured Yourdon-style models separate them into data processes, control processes and data stores, respectively. STATEMATE simulates local variables of statecharts by means of input and output flows accessible to the statechart. This means that in STATEMATE “local” data is really data present in communication channels between activities. STATEMATE models (as all other structured models) have separate activities where UML models only know of data manipulation local to an object. The absence of separate activities and the use of true local variables in the UML considerably simplifies our semantics w.r.t. the STATEMATE semantics. Secondly, to communicate information to a destination, objects in OO models must use the identifier(s) of the destination(s) whereas processes in structured models must use the identifiers of the communication channels. As a consequence, communication in OO models is point-to-point, whereas in structured models, it is broadcast. Finally, object-oriented models use the type-instance distinction, which is absent from structured models. This means that in our semantics, we can deal with dynamic creation and deletion of instances.

Table 2 summarises the differences between our semantics of UML statecharts and the semantics indicated informally in the OMG documentation [14]. The differences follow from the fact that we define a requirements-level semantics whereas the OMG defines an implementation-level semantics. The OMG UML 1.3 semantics makes the assump-

Table 2 Differences between requirements and implementation-level semantics

<i>requirements-level semantics</i>	<i>implementation-level semantics</i>
perfect technology	imperfect technology
input is event set	• input is queue
system reacts to all events in input	• system reacts to one event in input (the first)
event exists for one step	• event exists arbitrary long
instantaneous communication	• non-instantaneous communication
communication arrives always	communication may get lost, or arrive at wrong destination
one global clock	many local clocks
no clock drift	clocks may drift
action is instantaneous	• action takes time
unlimited concurrency	• limited concurrency, i.e., threads of control with interleaving per thread (one thread per active object)

tions marked by a bullet in table 2. Our semantics, like STATEMATE, makes all of the assumptions in the left-hand column.

Together, tables 1 and 2 factorise the differences between the STATEMATE and UML semantics of statecharts into two groups: the differences between structured and OO models and the difference between requirements-level and implementation-level semantics. On the other hand, Harel and Gery [6] state that the main difference is that UML statecharts use run-to-completion (RTC) whereas STATEMATE statecharts do not. RTC is one possible way to maintain atomicity of transitions at the implementation level. In an RTC semantics, an event can only be processed, if processing of the precious event input has completed (all triggered transitions have been completely taken). Sending an asynchronous message is considered to be completed when the message is sent; calling a synchronous operation is considered completed when the called operation is completed, and this requires maintaining a call stack. STATEMATE has no synchronous communication and hence no call stack. In our opinion, RTC is a minor difference compared to the differences identified by us.

We now briefly sketch some design choices we have made in defining a semantics. The following choices have to be made for any semantics of statecharts, regardless of semantic level or design paradigm (see also Von der Beeck [2]):

- We specify both a clock-synchronous and clock-asynchronous semantics [8]. In the clock-synchronous semantics, the system starts processing its input only at the tick of the clock. In the clock-

asynchronous semantics, the system starts processing its input as soon as it receives it. In the OMG semantics of UML, this issue is ignored.

- In synchronous communication, the caller must wait until the callee has finished processing the communication. In asynchronous communication, the caller continues without waiting for the receiver to finish processing the communication. STATEMATE only uses asynchronous communication. We follow the UML in defining a semantics for both. We show that in our requirements-level semantics, synchronous communication is only possible with a clock-asynchronous semantics.
- Like the UML [14], we allow no compound triggers, no negated trigger events, and only a single entry and single exit action.
- We allow for the case that the action expressions of a transition contain sequence, interleaving and parallel operators. The example action language that we use only contains the sequence operator, but our semantics can deal with the generalised case. Parallel actions that interfere lead nondeterministically to different possible states.
- Updates to variables are made at the end of a step. This way no inconsistent value can be read.
- We assume a given priority preorder on transitions, and do not commit ourselves to any particular one. Possible priority preorder are the STATEMATE one (a higher level transition has priority over a lower level one) and the UML one (a lower level transition has priority over a higher level one).
- In the UML, a distinction is made between active and passive objects. Active objects have computing resources, passive objects have not. At the requirements level, all objects have computing resources, so we only have active objects.

As most other statechart semantics, ours preserves causality (effect should not be contradictory to cause). We allow instantaneous states.

The following UML statechart constructs could be added to our requirements-level semantics without difficulty, but to simplify the exposition we omitted them:

- Deferred events. These can be simulated by regenerating an event as often as it is to be deferred.

- Parametrised events. These can be simulated by sets of parameterless events (one for every combination of parameter values). The resulting state space may be infinite (making model checking impossible).
- A taxonomy for events. This is an abbreviation mechanism that allows one to reduce the number of transitions in a statechart. It does not add expressive power.
- Activities (actions that take time). These can be simulated by instantaneous start and finish events.
- Dynamic choice points. Adding this UML construct would mean that a step is executed in two parts, separated by the dynamic choice point. Dynamic choice points can be simulated by adding an intermediary state.
- Synchronisation states. (These pseudo states are used to mimic a monitoring or semaphore construct.) We regard this to be an implementation-level construct. This construct can be simulated by a join (a transition with multiple source states) with a counter.
- History states. These can be simulated by (re)defining an entry function (see e.g. Damm *et al.* [3]).

Related work. Our LTS semantics is the UML analogon of the LTS semantics for STATEMATE statecharts defined by Damm *et al.* [3]. Due to the differences between UML and STATEMATE discussed above, our semantics is considerably simpler. Our semantics is based upon an earlier requirements-level UML semantics with simple state-transition diagrams given by Wieringa and Broersen [17]. All other approaches known to us are implementation-level semantics.

Latella *et al.* [10] use hierarchical automata to define a formal semantics for one UML statechart. The step semantics they define resembles our semantics of a step, but they do not define other aspects of behaviour (for example, communication).

Lilius and Porres [11] give an operational semantics for the original UML definition [14] by defining an algorithm for the execution of a step in pseudo-code. It is however unclear how to derive a Kripke structure from this algorithm. This is important since the authors have built a model checker based on their semantics. The authors do not define other aspects of behaviour (e.g. creation, deletion, communication).

Reggio *et al.* [16] give a semantics for one UML statechart belonging to an active object by defining a labelled transition system. They stick

as close as possible to the informal standard definition, eliminating inconsistencies, but they nevertheless have to make a number of choices mostly due to the implementation level. They do not define communication.

Harel and Kupferman [7] give an executable object model for a simplified version of a UML like system. They consider simple Mealy machines, whereas we consider statecharts. They focus on communication between different objects and on the possibilities of deadlock in a clock-asynchronous semantics. We, on the other hand, focus on communication at the requirements level where no deadlock is possible, and consider both a clock-synchronous and clock-asynchronous semantics.

In section 2, we summarise the necessary syntax definitions of UML statecharts. In section 3 we generalise the definition of a step so that it applies to structured and object-oriented design approaches and is independent from semantic level. In section 4 we then define the execution semantics of one statechart in terms of a labelled transition system. In section 5 we add creation and deletion action expressions and communication between statecharts. We end with a discussion and conclusions.

2. STATECHART SYNTAX

Most of the definitions in this section are adapted from Damm *et al.* [3] and Pnueli and Shalev [15] with minor changes to deal with final states and completion events in the UML. The paragraph on time expression is new. A statechart SC is a collection of state nodes, hierarchically related, and connected by edges. It is described by the tuple $SC = (Nodes, Edges, Events, Guards, Actions, Var)$, with $Nodes$ the set of state nodes, $Edges$ the set of edges, $Events$ the set of primitive event expressions, $Guards$ the set of guard expressions, $Actions$ the set of action expressions, and Var the set of local variables. We assume a set $Names$ to label the edges.

State nodes. The function $children : Nodes \rightarrow \mathbb{P}Nodes$ defines for each state the set of children of that state (its immediate substates). If $s \in children(s')$, we call s' the *parent* of s . If a state node has no children, it is *basic*. Otherwise, it is a *composite* state node. A composite state node is either an AND or an OR state node. When the system is in an OR state node, it is in one of its children. When the system is in an AND state node, it is in all of its children. The function $type : Nodes \rightarrow \{BASIC, AND, OR\}$ assigns to every state node its type.

Function $children$ defines a partial ordering on the state nodes:

$$s \leq s'$$

Reflexivity

$$\begin{aligned}
s &\leq s' && \text{if } s \in \text{children}(s') \\
s &\leq s' \wedge s' \leq s'' \Rightarrow s \leq s'' && \text{Transitivity}
\end{aligned}$$

If $s \leq s'$ we say that s' is the *ancestor* of s . If either $s \leq s'$ or $s' \leq s$, we say s and s' are *ancestrally related*.

There are some special state nodes. We assume a unique root state node, denoted *root*. We require *root* to be an OR state node. Another special state node is the final state node. If state node s is an OR state and one of its children is a final state node, that state node is denoted \mathbf{final}_s . A final state node must be basic, and no edges emanate from it. Obviously, $\mathbf{final}_s \leq \mathbf{final}_s$ and $\mathbf{final}_s \leq s$.

The partial function *default* : *Nodes* \rightarrow *Nodes* identifies for each OR state node s one of its children as the default state node, that is entered when s is entered.

For a set of state nodes $X \subseteq \text{Nodes}$, the least common ancestor $\text{lca}(X)$ is the state $l \in \text{Nodes}$ such that:

$$\begin{aligned}
\forall x \in X \cdot x &\leq l \\
\forall y \in \text{Nodes} \cdot (\forall x \in X \cdot x &\leq y) \Rightarrow l \leq y
\end{aligned}$$

Two state nodes s, s' are *orthogonal*, written $s \perp s'$, if s and s' are not ancestrally related and their *lca* is an AND state.

A set of state nodes $X \subseteq \text{Nodes}$ is *consistent* if, for every every pair $x, y \in X$, either x and y are ancestrally related or x and y are orthogonal. A *configuration* is a maximal consistent set of state nodes. Configurations are valid ‘states’ of the statechart.

In order to deal with UML completion events, we define predicate *stopped*(C, s) that is true if in configuration C (composite) state s is terminated:

$$\text{stopped}(C, s) = \begin{cases} \mathbf{false}, & \text{if } \text{type}(s) = \text{BASIC} \\ \mathbf{final}_s \in C, & \text{if } \text{type}(s) = \text{OR} \\ \forall x \in \text{children}(s) \cdot \text{stopped}(C, x), & \text{if } \text{type}(s) = \text{AND} \end{cases}$$

Edges. An edge is a relation from one set of state nodes (source) to another (target). An edge has a unique name. It is labelled with an event expression, a guard expression and an action expression. (Informally, when the event specified by the event expression occurs, and the guard expression evaluates to true, the action expression is evaluated.) Formally, *Edges* $\subseteq \text{Names} \times \mathbb{P}\text{Nodes} \times (\text{Events} \cup \emptyset) \times \text{Guards} \times \text{Actions} \times \mathbb{P}\text{Nodes}$. For every $t = (n, X, e, g, as, Y)$, let $\text{name}(t) = n$, $\text{source}(t) = X$, $\text{event}(t) = e$, $\text{guard}(t) = g$, $\text{action}(t) = as$ and $\text{target}(t) = Y$.

We require that $source(t)$ and $target(t)$ are non-empty consistent sets of state nodes.

For a given transition t whose source is composite state s , a special event that may be used to label t is \sqrt{s} . This we call the completion event. It occurs if and only if in a given configuration C , $stopped(C, s)$ becomes true.

The *scope* of a transition t is the most nested OR state that contains both $source(t)$ and $target(t)$. It is the most nested state of the statechart that is not exited or entered when the transition is taken.

Two transitions are consistent if they are equal or their scopes are orthogonal. A set of transitions T is called a *consistent* set if t_1 and t_2 are consistent, for every $t_1, t_2 \in T$. Denote by $consTrans(T)$ the set of all transitions that are consistent with every $t \in T$. Note that if T is consistent, then $T \subseteq consTrans(T)$.

Next, we assume a priority ordering \preceq on the transitions. We need an order to choose some transitions if we have a non-consistent set of transitions. We assume \preceq to be a total pre-order. We write $t \preceq t'$ to denote that t has higher or the same priority as t' .

Syntax of action expressions. The set of basic action expressions for a given statechart SC , denoted $Actions(SC)$, is defined as follows.

<code>v:=exp</code>	where $v \in Var(SC)$
<code>a;a'</code>	where $a, a' \in Actions(SC)$
<code>refid:=create(Class)</code>	where $refid \in Var(SC)$
<code>delete(id)</code>	
<code>send id.signal</code>	where <i>signal</i> is a signal reception of object <i>id</i>
<code>id.oper</code>	where <i>oper</i> is an operation of object <i>id</i>

This syntax may easily be extended to incorporate more advanced features. In section 5 we will explain the meaning of *Class* and *id*.

Time expressions. In the UML there are two kinds of time expressions: absolute and relative. An absolute time expression references the current time. It is specified by the UML $when(cond)$ construct (e.g. $when(time=12:00:00h)$). A relative time expression references the time relative to some occasion. It is specified by the UML $after(texp)$ construct (e.g. $after(10s)$, which is true 10 seconds after the source state has been entered), where $texp$ is of type \mathbb{N} .

Below, we will define a translation of $when$ and $after$ into basic guard expressions (i.e. not containing \wedge, \vee, \neg), that we call basic clock constraints. We assume that each basic clock constraint references a clock c of type real, and that the form of the constraint is $c \leq n$ or $c \geq n$ where $n \in \mathbb{N}$.

We translate an absolute time expression $\text{when}(\text{cond})$ into a guard expression that references a global clock (the current time) by treating cond as guard expression. We introduce special clock variable ct to represent the current time. At the beginning of the execution of the statechart, this clock is initialised with the current time.

We translate a relative time expression $\text{after}(\text{texp})$ as follows into a guard expression that references a local clock. We regard after as a special kind of event label that can only label transitions that have a single source. Let $x = \text{source}(t)$. We introduce a special clock variable tmr_x that denotes a local clock (a timer), that will be reset every time state node x is entered. The event $\text{after}(\text{texp})$ is replaced by guard expression $[\text{tmr}_x \geq \text{texp}]$.

Note that a basic clock constraint of the form $\text{tmr}_x \leq n$, where tmr_x is a timer and $n \in \mathbb{N}$, does not make sense. If state node x is entered, tmr_x is set to zero; hence, the basic clock constraint will be always true then.

In STATEMATE, scheduled actions may be used. A scheduled action, in STATEMATE syntax denoted $\text{sc}(a, \text{texp})$, means that action a must be done after delay texp . This can be specified in the UML using the after construct: $\text{after}(\text{texp})/a$.

3. GENERIC STATECHART STEP SEMANTICS

In this section and the next one we give a requirements-level semantics for one statechart. In this section we present the definition of a step. It is generic and independent from whether we follow a structured or object-oriented design approach. It is simpler than the STATEMATE definition of a step [8] [3] because there are neither negative nor compound events in the UML.

We define how to compute a step, given a configuration C and a set of input events I . This step must obey several constraints, such as maximality and consistency and priority. After Pnueli and Shalev [15], we define a step both declaratively and operationally. Note however that we do not define a superstep in this way, as they do, and that we take into account priorities, whereas they do not.

First, we define all enabled transitions $\text{En}(C, I)$:

$$\text{En}(C, I) = \{t \mid \text{source}(t) \subseteq C \wedge \text{event}(t) \in I \wedge \text{guard}_\sigma(t)\}$$

where $\text{guard}_\sigma(t)$ represents the evaluation of $\text{guard}(t)$ in valuation σ . (See section 4 for the definition of a valuation.)

We construct a step in an incremental way, by adding transitions that are consistent. To aid in this construction, we define function addToStep

Figure 1 *nextstep*(C, I): algorithm to compute steps

```

let  $T = \emptyset$ ;
while  $T \subset \text{addToStep}(T)$  do
  choice  $t \in \text{high}(\text{addToStep}(T) - T)$ ;
  let  $T = T \cup \{t\}$ 
endwhile
return  $T$ 

```

as follows. Suppose we have decided to take a set of transitions T . The function $\text{addToStep}(T)$ gives us the transitions that are enabled and are consistent with T . The function $\text{addToStep}(T)$ is defined by:

$$\text{addToStep}(T) = \text{En}(C, I) \cap \text{consTrans}(T)$$

A step T^* must satisfy the following constraints:

- 1 *Enabledness* All transitions of the step must be enabled.
- 2 *Consistency* All transitions taken in T^* are consistent with each other, i.e., each of them could be added.
- 3 *Maximality* The step is maximal, i.e., every transition that could be added is already part of the step.
- 4 *Priority* If an enabled transition is not in the step, then there must a transition in the step that is inconsistent and has higher or the same priority.

We formalise these constraints as follows:

$$\begin{array}{ll}
T^* \subseteq \text{En}(C, I) & \text{Enabledness} \\
T^* \subseteq \text{consTrans}(T^*) & \text{Consistency} \\
\text{addToStep}(T^*) \subseteq T^* & \text{Maximality} \\
\forall t \in \text{En}(C, I) - T^* \exists t' \in T^* \cdot t \not\subseteq \text{consTrans}(\{t'\}) \wedge t' \preceq t & \text{Priority}
\end{array}$$

Note that the first three constraints are equivalent to the fixpoint equation $T^* = \text{addToStep}(T^*)$.

Next, we give a nondeterministic algorithm $\text{nextstep}(C, I)$ to compute the maximal, nonconflicting step given a configuration C and a set of input events I (see Fig. 1). In the definition of the algorithm, we use a set $\text{high}(T)$. This set $\text{high}(T) \subseteq T$ contains the transitions of T with the highest priority:

$$\text{high}(T) = \{t \in T \mid \forall t' \in T \cdot t \preceq t'\}$$

Now that we have defined a step in two different ways, we show the consistency of both approaches. We call a step T^* constructible if and only if T^* can be constructed via procedure $nextstep(C, I)$.

Proposition 1. T^* satisfies the constraints 1-4 iff T^* is constructible.

Proof. Omitted due to space limitations. See [4]. \square

4. STATECHART LTS SEMANTICS

We now present two different ways of executing a step. In the *clock-synchronous* semantics, a step is executed when the clock ticks. In the *clock-asynchronous* semantics, a step is executed immediately when new events arrive. Then the system becomes unstable and reacts infinitely fast to become stable again. These two semantics are borrowed from the STATEMATE semantics. The major differences with STATEMATE semantics are the different treatment of termination in the UML and the absence of a separate activity model in the UML. Besides, our definition of real time is different from the STATEMATE definition.

4.1. CLOCKED LABELLED KRIPKE STRUCTURE

We assume a typed set of variables Var and a typed data domain \mathcal{D} . A *valuation function* $\sigma : Var \rightarrow \mathcal{D}$ assigns to every variable a value. The set of valuation functions defined on Var is denoted $\Sigma(Var)$. Every state of the clocked labelled Kripke structure is a valuation σ of the variables Var . From now on, we will use the term ‘valuation’ instead of ‘state of a clocked labelled Kripke structure’ to avoid confusion with the state nodes of statecharts. Valuations are connected by transitions labelled with actions out of set Act .

A clocked labelled Kripke structure (CLKS) is a tuple $(Var, Act, \rightarrow, ci, \sigma_0)$ with:

- Var the set of variables,
- Act the set of actions,
- $\rightarrow \subseteq \Sigma(Var) \times \mathbb{P}Act \times \Sigma(Var)$ the transition relation,
- $ci : \sigma \rightarrow CC$ a function that assigns a conjunction of basic clock constraints to every valuation, the clock invariant,
- $\sigma_0 \in \Sigma(Var)$ the initial valuation.

Set CC denotes all possible conjunctions of all basic clock constraints. Instead of $(\sigma, A, \sigma') \in \rightarrow$, we write $\sigma \xrightarrow{A} \sigma'$. The clock invariant is a

conjunction of basic clock constraints of the form $c \leq n$. Given configuration $\sigma(cfg)$, for every transition t with $source(t) \subseteq \sigma(cfg)$, that has a guard expression containing a clock constraint of the form $c \geq n$, the clock invariant ci_σ contains a basic clock constraint $c \leq n$. Below, we will see that the clock-invariant helps in raising a temporal event (see the definition of \rightarrow_{time}). This is necessary for the clock-asynchronous semantics. It is not necessary for the clock-synchronous semantics, since at every tick of the clock input is processed, but it helps in clarifying the clock-synchronous semantics.

For readers familiar with timed automata [1] or clocked transition systems [12], note that we do not have defined clock constraints on transitions of the CLKS. Instead, clock constraints are present as guard expressions in the statechart! We could have defined clock constraints on the transitions of the CLKS, but this would have complicated our semantics considerably without adding new information to the semantic structure. The set of clocks to be reset is implicit; see the definition of \rightarrow_{copy} below.

When interpreting a statechart in an LTS, the set of variables Var comprises the variables of a statechart SC (denoted $Var(SC)$), and a copy of these variables, denoted by priming the original variables. Updates are made to the primed variables, to ensure that actions can only read the values variables had at the beginning of a step.

Special elements of Var are the variables cfg and inp denoting resp. the current configuration of the statechart and the input of the statechart, represented as a set of events. Next, we have a clock variable ct that denotes the current time, and a set of clock variables tmr that contains all timers tmr_x , where x is the source of a transition that has a guard expression referencing local clock tmr_x (see section 2).

We only consider Kripke structures that satisfy the constraint that all relevant completion events are in the input:

$$\{\checkmark_s \mid stopped(\sigma(cfg), s) \wedge s \in \sigma(cfg)\} \subseteq \sigma(inp)$$

We let set Act be the union of the sets $Events$, $Actions$, and $Guards$. We assume a function $act : Edges \rightarrow Act$ that, given a set of edges, returns the set of actions that are in the set. This function will be used to label the transitions in the CLKS with actions.

Next, we sketch briefly how action expressions are evaluated. If a is an action expression and σ a valuation, then $\llbracket a \rrbracket \sigma$ is the valuation resulting from the evaluation of expression a . Updates are made to the primed versions of variables only. For example:

$$\begin{aligned} \llbracket x := exp \rrbracket \sigma &= \sigma[x' / \llbracket exp \rrbracket \sigma] \\ \llbracket action; actionsequence \rrbracket \sigma &= \llbracket action \rrbracket \sigma; \llbracket actionsequence \rrbracket \sigma \end{aligned}$$

The meaning of the x/val construct on valuations is defined as follows:

$$\sigma[x/val](y) = \begin{cases} \sigma(y), & \text{if } y \neq x \\ val, & \text{if } y = x \end{cases}$$

Finally, the initial valuation σ_0 must satisfy $\sigma_0(cfg) = dcomp(\{root\})$, $\sigma_0(input) = \emptyset$, $\forall t \in tmr \cdot \sigma_0(t) = 0$, and ct and the local variables must be given an appropriate value. Function $dcomp$ computes the default completion of a set of state nodes (see e.g. [15], [3] for a definition).

4.2. BASIC CLOCK-SYNCHRONOUS EXECUTION

The \rightarrow relation of the Kripke structure is the union of three relations \rightarrow_{time} , \rightarrow_{change} , and \rightarrow_{step} , defined below. Not every sequence of transitions out of these three sets satisfies the clock-synchronous semantics. Sequences in the clock-synchronous semantics are cycles, where a cycle describes the execution of the CLKS during one time-unit (from one clock tick to another). A cycle is a non-empty sequence of time steps (which represent the elapsing of time), possibly interleaved with change steps (which represent changes in the environment), followed by the taking of a step at the tick of the clock:

$$\rightarrow_{cycle} = \xrightarrow{\delta_0}_{time}; \rightarrow_{change}; \xrightarrow{\delta_1}_{time}; \rightarrow_{change} \dots \xrightarrow{\delta_j}_{time}; \rightarrow_{step}$$

where $\sum_{i=0}^j \delta_i = 1$. Relation \rightarrow_{time} corresponds to the elapsing of time. Relations \rightarrow_{change} represents a change in the environment. There are three possible changes: either an external event occurs (\rightarrow_{event}), a value change occurs (\rightarrow_{value}), or a temporal event occurs ($\rightarrow_{time-out}$).

$$\begin{aligned} \sigma \xrightarrow{\delta}_{time} \sigma' &\Leftrightarrow \sigma' = \sigma[ct/(\sigma(ct) + \delta), \\ &\quad \&_{tmr_x | tmr_x \in tmr} tmr_x / (\sigma(tmr_x) + \delta)] \\ &\wedge \forall \varepsilon \in [0, \delta] \cdot ci_{\sigma''} \quad \text{where} \\ &\quad \sigma'' = \sigma[\&_{tmr_x | tmr_x \in tmr} tmr_x / (\sigma(tmr_x) + \varepsilon)] \\ \rightarrow_{change} &= \rightarrow_{event} \cup \rightarrow_{value} \cup \rightarrow_{time-out} \\ \sigma \rightarrow_{event} \sigma' &\Leftrightarrow \exists E \subseteq Events \cdot E \neq \emptyset \\ &\quad \wedge \sigma' = \sigma[inp/(\sigma(inp) \cup E)] \\ \sigma \rightarrow_{value} \sigma' &\Leftrightarrow \sigma \neq \sigma' \\ &\quad \wedge \sigma(cfg) = \sigma'(cfg) \\ &\quad \wedge \sigma(inp) = \sigma'(inp) \\ &\quad \wedge \forall tmr_x \in tmr \cdot \sigma(tmr_x) = \sigma'(tmr_x) \\ \sigma \rightarrow_{time-out} \sigma' &\Leftrightarrow \sigma = \sigma' \end{aligned}$$

$$\wedge \nexists \sigma'' \cdot \sigma \rightarrow_{time} \sigma''$$

Relation $\xrightarrow{\delta}_{time}$ represents the increase of the global clock ct and all timers tmr_x by δ , provided that during this period all clock invariants hold. The ampersand symbol denotes that all updates to all timers are made simultaneously:

$$\sigma[\&_{id \in ID} x_{id}/val_{id}] = \sigma[x_{firstid}/val_{firstid} \dots x_{lastid}/val_{lastid}]$$

where $ID = \{firstid \dots lastid\}$. Relation \rightarrow_{event} extends the current input set nondeterministically with a set E without changing the system state. Relation \rightarrow_{value} indicates a value change of one or more local variables. Relation $\rightarrow_{time-out}$ represents that a temporal event (time-out) occurs iff a time step is not possible. Next, we define the \rightarrow_{step} relation, corresponding to the execution of a step.

$$\begin{aligned} \xrightarrow{A}_{step} &= \xrightarrow{A}_{nextstep}; \rightarrow_{copy} \\ \sigma \xrightarrow{A}_{nextstep} \sigma' &\Leftrightarrow A = act(T^*) \\ &\wedge \sigma' = \sigma''[cfg/nextconfig(\sigma(cfg), T^*), inp/\emptyset] \\ &\quad \text{where } isStep(T^*) \text{ and } \sigma'' = \llbracket actions(T^*) \rrbracket \sigma. \\ \sigma \rightarrow_{copy} \sigma' &\Leftrightarrow \sigma' = \sigma[\&_{v \in Var(SC)} v/\sigma(v) \\ &\quad \&_{tmr_x | tmr_x \in tmr \wedge x \in \sigma(cfg)} tmr_x/0] \end{aligned}$$

Relation $\xrightarrow{A}_{nextstep}$ relates valuation σ with σ' iff given input events $\sigma(inp)$ in configuration $\sigma(cfg)$ step T^* , satisfying $isStep(T^*)$ is taken, and the actions of T^* are evaluated, such that valuation σ'' is reached. Next, the configuration and inputs are updated, the input is removed, and valuation σ' is reached. (See e.g. [3] for a precise definition of function $nextconfig$.) Predicate $isStep(T^*)$ is true iff step T^* satisfies the consistency, maximality and priority constraints, given configuration $\sigma(cfg)$ and set of inputs $\sigma(inp)$. Whether completion transitions have priority over non-completion transitions, depends on the definition of the priority relation \preceq . Function $actions(st)$, given a set of transitions, returns a sequence of their actions, respecting the order of the action within each transition, but nondeterministically choosing an order between the transitions themselves:

$$actions(T) = action(t_1); \dots; action(t_n) \Leftrightarrow \{t_1, \dots, t_n\} = T$$

Relation \rightarrow_{copy} represents that all unprimed variables are updated with the values of their primed counter parts and all relevant timers are reset.

4.3. BASIC CLOCK-ASYNCHRONOUS EXECUTION

The \rightarrow relation of the CLKS is a union of the \rightarrow_{time} , \rightarrow_{change} and $\rightarrow_{superstep}$ relations. Like in the clock synchronous semantics, an execution sequence is a sequence of cycles. Now, however, a cycle describes the execution of the CLKS from the end of one reaction until the end of the execution of the next one. A cycle is a sequence of time steps, followed by a change step, immediately followed by the reaction of the system, a superstep. A superstep is a chain of steps. The fact that a chain of steps is taken, rather than a single step, is due to the perfect technology assumption (reactions happen infinitely fast) and the assumption of the clock-asynchronous semantics that a change in the environment, including one caused by the system itself, is immediately reacted upon. This implies that all completion events are immediately reacted upon; hence in all valuations the configuration of the statechart is stable (does not contain a final state node). Another effect of this definition is that a superstep may be infinite. In that case, we say the superstep diverges.

$$\begin{aligned} \frac{A \rightarrow_{cycle}}{\sigma \xrightarrow{A} superstep \sigma'} &= \rightarrow_{time}; \rightarrow_{change}; \frac{A \rightarrow_{superstep}}{\sigma \xrightarrow{A_1} step \dots \xrightarrow{A_n} step \sigma'} \end{aligned}$$

where $A = A_1 \cup \dots \cup A_n$ and $stable(\sigma'(cfg), \sigma'(inp))$ must hold:

$$stable(C, I) \Leftrightarrow (\forall s \in C \cdot stopped(C, s)) \wedge I = \emptyset \wedge En(C, I) = \emptyset$$

The superstep stops if the current configuration is stable and there are no more inputs and no more enabled transitions. The other definitions are as before.

5. COMMUNICATING STATECHARTS

We now add creation and deletion of objects, absent from STATEMATE semantics, and synchronous and asynchronous communication. STATEMATE only uses asynchronous communication. Our semantics leans heavily on the use of object identifiers.

Assume a set of Classes *Class* and a set of object identifiers *OID*. Every class *c* has a statechart definition associated, whose root is identified as *c.root*. Now that we have multiple objects instead of one, we index the special variables *cfg*, *inp* and *tmr* with the object id's. So *cfg[id]* denotes the configuration *cfg* of object *id*. We assume that all variables used by objects are unique, including clock variables. (Variables can be made unique by simply putting the object id in front.) We denote the local variables of an object *id* with *Var(id)*. As before, every variable

$v \in \text{Var}(id)$ has a primed version, but now, in addition, every set of input events $\text{inp}[id]$ has also its primed version $\text{inp}'[id]$, because we define communication between statecharts.

We only consider CLKs for which it holds that for every existing object, all relevant completion events are in its input. The formalisation of this, and the definition and formalisation of the initial valuation of the CLKs are straightforward generalisations of the single-object case. We refer the interested reader to [4].

5.1. CREATION AND DELETION

The meaning of action expression $\text{refid} := \text{create}(\text{Class})$ is that an object of class Class is created with a new identity whose value is assigned to variable refid . Action expression $\text{destroy}(id)$ means that object id is destroyed.

In order to give a formal semantics for creation and deletion we make use of two predicates. Predicate $\text{Exists}(id)$ is true iff an object with identifier id exists. Predicate $\text{Used}(id)$ is true iff there exists or existed an object with identifier id . The obvious constraint holds that $\text{Exists}(id) \Rightarrow \text{Used}(id)$.

$$\begin{aligned} \llbracket \text{refid} := \text{create}(\text{Class}) \rrbracket \sigma &= \sigma[\text{cfg}[id]/\text{dcomp}(\{\text{Class.root}\}), \\ &\quad \text{inp}[id]/\emptyset, \\ &\quad \text{refid}/id, \\ &\quad \text{Exists}(id)/\text{true}, \\ &\quad \text{Used}(id)/\text{true}] \quad \text{if } \neg \text{Used}_\sigma(id) \\ \llbracket \text{destroy}(id) \rrbracket \sigma &= \sigma[\text{Exists}(id)/\text{false}] \quad \text{if } \text{Exists}_\sigma(id) \end{aligned}$$

A new object id can be created only if id is not used before. If object id can be created, then it is initialised by setting its configuration and set of input events, and updating predicates Used and Exists .

5.2. COMMUNICATION

Clock-synchronous model. In the clock-synchronous model, we only allow asynchronous communication (the **send** action). The reason is that in synchronous communication, the caller must wait until the callee returns. In the clock-synchronous model, the callee is performing its own step when it receives the call and can respond to the call only in the next step at the next tick of the clock. Since in our requirements-level semantics, a transition must be taken in zero time, the caller cannot

First, observe that in the clock-asynchronous model both asynchronous and synchronous communication is allowed. Synchronous communication is allowed because each object instantly reacts to the events it receives and hence always is ready to synchronise with another object. We assume that only a single object is active during a single step. We can justify this by our other assumption that single steps and supersteps do not take time.

A step, executed by one object id , is defined as follows. The definition formalises RTC semantics.

$$\begin{aligned}
\sigma \xrightarrow{A, id} step \sigma' &= \sigma \xrightarrow{A, id} nextstep; \xrightarrow{id} copy \\
\sigma \xrightarrow{A, id} nextstep \sigma' &\Leftrightarrow A = act(T^*) \\
&\quad \wedge \sigma' = \sigma''[cfg[id]/nextconfig(\sigma(cfg[id]), T^*)] \\
&\quad \text{where } isStep(T^*) \text{ and } \sigma'' = \llbracket actions(T^*) \rrbracket \sigma \\
\sigma \xrightarrow{id} copy \sigma' &\Leftrightarrow \sigma' = \sigma[\&_{v \in Var(id)} v/\sigma(v'), \\
&\quad inp[id]/\sigma(inp'[id]), \\
&\quad inp'[id]/\emptyset]
\end{aligned}$$

The above definition is a straightforward adaptation of the previous definitions of $\rightarrow step$, $\rightarrow nextstep$ and $\rightarrow copy$. The change w.r.t. the clock-synchronous model of this section is that we now define a step for one object only. A superstep S of a finite set of objects O will be a non-deterministic interleaving of supersteps of the each object in O . The next definition says that in each single step of S we nondeterministically choose a step of some object in O .

$$\begin{aligned}
\sigma \xrightarrow{A} singlestep \sigma' &\Leftrightarrow \exists id \in OID \mid Exists_{\sigma}(id) \cdot \\
&\quad \wedge \sigma \xrightarrow{A, id} step \sigma' \\
\sigma \xrightarrow{A} superstep \sigma' &\Leftrightarrow \sigma \xrightarrow{A_1} singlestep \dots \xrightarrow{A_n} singlestep \sigma' \\
&\quad \wedge \forall id \in OID \mid Exists_{\sigma'}(id) \cdot \\
&\quad \quad stable(\sigma'(cfg[id]), \sigma'(inp[id]))
\end{aligned}$$

where $A = A_1 \cup \dots \cup A_n$.

What remains is to define the meaning of synchronous calls as part of the execution of $actions(T^*)$. If an object calls operation $oper$ of object id , then first of all id should exist, and second, id should process the call. Since id itself may have nonempty input, this old input should be remembered and placed back after id has finished executing. Together with $\rightarrow step$, this formalises RTC semantics, because it says that a call action is executed only when the called action is executed.

$$\begin{aligned}
[[id.oper]]\sigma = \sigma' \text{ if } & \text{Exists}_\sigma(id) \\
& \wedge \sigma_1 \xrightarrow{A, id} \text{step}\sigma_2 \\
& \wedge \sigma' = \sigma_2[inp[id]/(\sigma_2[inp[id]] \cup \sigma[inp[id]])] \\
& \text{where } \sigma_1 = \sigma[inp[id]/oper]
\end{aligned}$$

6. CONCLUSION

This paper presents a formal real-time requirements-level semantics for UML statecharts. The step relation is specified both declaratively and operationally and is parametrised with respect to the priority definition. The most important aspect of the semantics is the definition of communication between several co-existing and cooperating objects. The semantics supports dynamic instantiation of objects. As explained in the introduction, our semantics makes it possible to identify the differences of UML statecharts with STATEMATE statecharts.

Because it is defined in terms of a transition system, our semantics can be used to define an execution semantics for UML models at the requirements level. In addition, it is the semantic basis for model checking UML requirements models. First results in model checking are reported elsewhere [9].

Future work includes, next to model checking, the study of inheritance in OO systems by applying techniques from concurrency (similar to e.g. [7]).

Acknowledgments

The authors would like to thank Maarten Fokkinga and David Jansen for commenting on previous parts of this paper.

References

- [1] R. Alur and D.L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
- [2] Beeck, M. von der. A comparison of statecharts variants. In H. Langmaack, W.-P. de Roever, and J. Vytupil, editors, *Formal Techniques in Real-Time and Fault-Tolerant Systems*, Lecture Notes in Computer Science 863, pages 128–148. Springer-Verlag, 1994.
- [3] W. Damm, B. Josko, H. Hungar, and A. Pnueli. A compositional real-time semantics of STATEMATE designs. In W.-P. de Roever, H. Langmaack, and A. Pnueli, editors, *Proc. COMPOS 97*, Lecture Notes in Computer Science 1536. Springer-Verlag, 1997.
- [4] R. Eshuis and R. Wieringa. Requirements-level semantics for UML statecharts. Technical Report TR-CTIT-00-07, University of

Twente, 2000.

- [5] D. Harel. Statecharts: A visual formulation for complex systems. *Science of Computer Programming*, 8(3):231–274, 1987.
- [6] D. Harel and E. Gery. Executable object modeling with statecharts. *IEEE Computer*, 30(7):31–42, 1997.
- [7] D. Harel and O. Kupferman. On the behavioral inheritance of state-based objects. Technical Report MCS99-12, Weizmann Institute of Science, 1999.
- [8] D. Harel and A. Naamad. The STATEMATE Semantics of Statecharts. *ACM Transactions on Software Engineering and Methodology*, 5(4):293–333, 1996.
- [9] D. N. Jansen and R. J. Wieringa. Extending CTL with actions and real-time. 2000. Submitted.
- [10] D. Latella, I. Majzik, and M. Massink. Towards a formal operational semantics of UML statechart diagrams. In *Proc. FMOODS'99, IFIP TC6/WG6.1*. Kluwer, 1999.
- [11] J. Lilius and I. Porres Paltor. Formalising UML state machines for model checking. In R. France and B. Rumpe, editors, *Proc. UML'99*, Lecture Notes in Computer Science 1723. Springer-Verlag, 1999.
- [12] Z. Manna and A. Pnueli. Clocked transition systems. In A. Pnueli and H. Lin, editors, *Logic and Software Engineering*. World Scientific, 1996.
- [13] S. McMenamin and J. Palmer. *Essential Systems Analysis*. Yourdon Press, New York, New York, 1984.
- [14] OMG. *Unified Modeling Language version 1.3*. OMG, July 1999.
- [15] A. Pnueli and M. Shalev. What is in a step: On the semantics of statecharts. In T. Ito and A. R. Meyer, editors, *Theoretical Aspects of Computer Software*, Lecture Notes in Computer Science 526, pages 244–265. Springer-Verlag, 1991.
- [16] G. Reggio, E. Astesiano, C. Choppy, and H. Hussmann. Analysing UML active classes and associated state machines – a lightweight formal approach. In T.S.E. Maibaum, editor, *Proc. FASE 2000*, Lecture Notes in Computer Science 1783. Springer-Verlag, 2000.
- [17] R. Wieringa and J. Broersen. A minimal transition system semantics for lightweight class- and behavior diagrams. In M. Broy, D. Coleman, T.S.E. Maibaum, and B. Rumpe, editors, *Proc. PSMT'98*. Technische Universität München, TUM-I9803, 1998.