

FORMAL SEMANTICS FOR INTERACTING UML SUBSYSTEMS

Jan Jürjens*

Department of Informatics, TU Munich, Germany

Computing Laboratory, University of Oxford, GB

<http://www.jurjens.de/jan> – juerjens@in.tum.de

Abstract So far most work on formal semantics for the Unified Modeling Language (UML) has concentrated on single diagrams. To provide a formal foundation for complete system specifications, one needs to put models for the different diagrams into context. We provide a formal semantics for UML subsystems that incorporates a formal semantics of the diagrams contained in a subsystem. It provides message-passing between objects or components specified in different diagrams, including a dispatching mechanism for events, and the handling of actions. It enables one to compose subsystems from sets of subsystems and allows them to interact by passing messages.

We give consistency conditions for the diagrams in a subsystem and define a notion of behavioural equivalence and two kinds of refinement for UML subsystems.

Keywords: UML, formal semantics, refinement, Abstract State Machines.

1. Introduction

The Unified Modeling Language (UML) (Rumbaugh et al., 1999) is the de-facto industry standard for specifying object-oriented software systems (for an introduction cf. (Stevens and Pooley, 2000)). Even though rather precisely defined compared with other modelling languages, its semantics is given only in prose form (UML, 2001), leaving room for ambiguities (a problem especially when providing tool support or trying to establish behavioural properties of UML specifications). Thus we need a mathematically precise semantics for UML.

There has been a substantial amount of work towards providing a formal semantics for UML diagrams (including (Evans et al., 1999; Börger et al.,

*Supported by the Studienstiftung des deutschen Volkes when carrying out this research.

2000; Stevens, 2001)). However, most work only provides models for single UML diagrams in isolation. When trying to give a precise mathematical meaning to whole UML specifications, one needs to be able to combine the formal models for the different kinds of diagrams.

Joint formal execution. In this paper we describe some results on how to formally model UML diagrams in context. We provide a formal semantics for UML subsystems that incorporates the formal semantics of the diagrams contained in the subsystem (here we consider statechart and activity diagrams; a treatment of the remaining diagrams is given in (Jürjens, 2002)). Specifically, our semantics

- models actions and internal activities explicitly (rather than treating them as atomic given events), in particular the operations and their parameters employed in them,
- provides message-passing between objects or components specified in different diagrams, including a dispatching mechanism for events and the handling of actions, and thus
- allows whole specification documents to be based on a formal foundation.

In particular, we show how to *compose* subsystems by including them into other subsystems.

Thus it prepares the ground for further work that may

- provide tool-support based on this precise semantics, in particular allowing complete specifications to be simulated, and
- provide the possibility of complete executable UML specifications.

Statechart, activity diagrams, subsystems. The semantics for statecharts presented here is based on (Börger et al., 2000), which however had to be extended significantly to incorporate the features mentioned above. Our formal semantics for subsystems and their interactions seems to be the first given in the published literature.

Consistency. Furthermore we give some conditions for *consistency*, relating different diagram kinds in a UML system specification, which is an issue that seems to attract increasing attention.

Equivalence. We define a notion of behavioural equivalence between UML subsystems. This can be used e.g. to verify consistency of two subsystem specifications that are supposed to describe the same behaviour, e.g. one of which uses statecharts to specify object behaviour, and the other sequence diagrams (the treatment of sequence diagrams can be found in (Jürjens, 2002)).

Refinement. In UML, *refinement* denotes a certain kind of dependency relation between model elements (UML, 2001, 2-18). There is no constraint on the semantic relationship between the model elements (also in the heuristics for state machine refinement at (UML, 2001, 2-177)).

When trying to establish system properties, behavioural conformance of refinement can lead to great savings (properties may be easier to establish at a more abstract level; preservation by refinement means that this is in fact sufficient). We define a kind of refinement (inspired by (Broy and Stølen, 2001)) that provides full behavioural conformance.

UML for security. One of the motivations for this work is to provide a foundation for using UML for developing security-critical systems (Jürjens, 2001-b), where a mathematically precise modeling is indispensable. The semantics presented here has been extended to include possible adversary behaviour to reason about security requirements (Jürjens, 2002). In particular, the preservation results in (Jürjens, 2001-a) extend to the refinement presented here.

For readability, we consider simplified UML diagrams. Our approach works just as well without the simplifications. More details can be found in (Jürjens, 2002).

Outline. In Section 2 we give basic definitions of Abstract State Machines needed for our semantics. We provide a sketch of the semantics of statechart diagrams and activity diagrams in Section 3, in order to be able to show how these fit together in the context of UML subsystems, the semantics of which is presented in Section 3.3. We define two kinds of refinement for UML subsystems in Section 4 and give an example in Section 5. We end with pointers to related work, a conclusion and indication of future work.

2. Abstract State Machines

ASMs (Gurevich, 1995) provide a rather flexible framework for formal modelling. They have been used e.g. to give a formal semantics for a non-trivial part of Java (Stärk et al., 2001).

We collect some central concepts. A *vocabulary* is a finite collection of function names, each of a fixed arity. We assume a set of *variables*. A *state* A of vocabulary $\text{Voc}(A)$ is a non-empty set X containing distinct elements *true*, *false*, and *undef* together with interpretations of the function names in $\text{Voc}(A)$ on X . Relations and sets are viewed as functions taking values in $\{\text{true}, \text{false}\}$. An ASM is executed by *updating* its state (e.g. changing the interpretations of the names in the vocabulary) iteratively by applying *update rules* some of which are given in the following.

Update The *update rule* $f(\bar{s}) := t$ of the ASM A updates f at the tuple \bar{s} to map to the element t .

Conditional For a Boolean term g rules R, S , the rule **if** g **then** R **else** S executes R if g holds, otherwise S .

Blocks The rule **do – in – parallel** R_1, \dots, R_k **enddo** executes the rules R_1, \dots, R_k simultaneously if they are mutually consistent: for any two update rules $f(\bar{s}) := t$ and $f(\bar{s}) := t'$, we have $t = t'$. Otherwise, execution stops.

Sequential Composition For rules R, S , the rule **seq** R, S **endseq** executes R and S sequentially.

Do-forall For a variable v , a Boolean term $g(v)$, and a rule $R(v)$, the rule **forall** v **with** $g(v)$ **do** $R(v)$ executes $R(a)$ for all a such that $g(a)$ holds, if they are mutually consistent. Otherwise, execution stops.

Loop through list For a variable v , a list X , and a rule $R(x)$, the rule **loop** v **through list** X $R(v)$ executes $R(x)$ iteratively for all $x \in X$.

An *abstract state machine* consists of a set of states and an update rule. It is executed by iteratively firing the update rule.

For multi-sets, we write $\{\!\{ \}$ instead of the usual brackets. For two multi-sets M and N , $M \uplus N$ denotes their union and $M \setminus N$ the subtraction of N from M . For a multi-set M and a set X , we write $M \setminus X$ for the multi-set of those elements in M (preserving their cardinalities) that are also elements of X .

3. Formal Semantics for a fragment of UML

The set **MsgNm** of *message names* consists of finite sequences of names $n_1.n_2.\dots.n_k$ where n_1, \dots, n_{k-2} are names of UML subsystems (to be defined below), n_{k-1} is a name of an object, and n_k is the local name of the message. The idea is that a message $n_1.n_2.\dots.n_k$ will be delivered as the message with name n_k to the object with name n_{k-1} which is part of the (iteratively nested) sequence of subsystems n_{k-2}, \dots, n_1 . Messages in **MsgNm** can be operations, signals, and return messages. For each operation op there is a corresponding return signal $\text{return}(op)$, assumed to be given explicitly. We write **Events** for the set of *events* of the form $op(exp_1, \dots, exp_n)$ with $op \in \text{MsgNm}$ and $exp_1, \dots, exp_n \in \text{Exp}$ (for a set **Exp** of *expressions* which includes a set **Var** $\subseteq \text{Exp}$ of variables). We define $\text{msgname}(m) \stackrel{\text{def}}{=} \text{msg}$ to be the name of the message $m = \text{msg}(exp_1, \dots, exp_n)$. In our model, every object O has associated multi-sets $\text{inQueue}(O)$ and $\text{outQueue}(O)$. We model sending a message msg from an object S to an object R as follows:

- (1) The object S puts $R.\text{msg}$ into its multi-set $\text{outQueue}(S)$.

- (2) A scheduler distributes the messages from out-queues to the intended in-queues (removing the head); in particular, $R.msg$ is removed from $outQueue(S)$ and msg added to $inQueue(R)$.
- (3) The object R removes msg from its in-queue and processes its content.

In the case of operation calls, we also need to keep track of the sender to allow sending return signals. Modelling communication is thus very flexible; we can adjust the scheduler according to the underlying communication layer (and take account e.g. of security (Jürjens, 2001-b) or performance issues).

In the following sections, we define the abstract syntax of various UML diagrams using mathematical notation, and then give a precise semantics of the modelled system behaviour using ASMs. For space reasons, we can only present simplified versions of the diagrams; more details are in (Jürjens, 2002).

3.1. Statechart Diagrams

Abstract syntax. A statechart diagram $D = (Obj_D, Cls_D, Sts_D, Ini_D, Trs_D)$ is given by an object name Obj_D , a class name Cls_D , a set of *states* Sts_D , an *initial state* Ini_D , and a set of transitions Trs_D . Sts_D is a set of tuples $S = (nm(S), ety(S), init(S), sta(S), int(S), ext(S))$ where

- $nm(S)$ is a string of characters called the *name* of the state,
- $ety(S) \in \mathbf{Action}$ is called the *entry action*,
- $init(S) \in Sts_D \cup \{undef\}$ is the *initial substate* of S ,
- $sta(S) \subseteq Sts_D$ is the set of substates of S ,
- $int(S)$ is called the *internal activity* (or *do-activity*), and
- $ext(S) \in \mathbf{Action}$ is the exit action.

Here we write **Action** for the set of actions of the following forms:

Call/send action: $call(msg(a_1, \dots, a_n))$ resp. $send(msg(a_1, \dots, a_n))$ for $msg \in \mathbf{MsgNm}$ and $a_i \in \mathbf{Exp}$.

Assignment: $att := exp$ for an attribute att and $exp \in \mathbf{Exp}$.

We assume that for every internal activity $actv$ there is an ASM rule **ActvRule**($actv$).

Trs_D is a set of tuples $t = (src(t), evt(t), grd(t), act(t), tgt(t))$ where

- $src(t) \in Sts_D$ is the *source state* of t ,
- $evt(t) \in \mathbf{Events}$ is the triggering event of t ,
- $grd(t)$ is a Boolean expression called the *guard* of t ,
- $act(t) \in \mathbf{Action}$ is an action (to be performed when firing t),

- $\text{tgt}(t) \in \text{Sts}_D$ is the *target state* of t .

$\text{evt}(t)$ must be of the form $\text{op}(\text{exp}_1, \dots, \text{exp}_n)$ with mutually distinct $\text{exp}_1, \dots, \text{exp}_n \in \mathbf{Var}$. We assume a special event $\text{CompEv} \in \mathbf{Events}$ (without parameters) as in (Börger et al., 2000). If $\text{intern}(t) = \text{true}$ then t is called an *internal* transition, otherwise it is called *external*.

Behavioural semantics. We give a formal semantics of statechart diagrams using ASMs. It is based on (Börger et al., 2000), which however had to be extended significantly to incorporate explicit modelling of actions and internal activities and message-passing between objects or components in different diagrams. We fix a statechart diagram D modelling an object $O \stackrel{\text{def}}{=} \text{Obj}_D$ and give its behavioural semantics as an ASM $\llbracket D \rrbracket^{SC}$.

The vocabulary of $\llbracket D \rrbracket^{SC}$ consists of the following names:

- the set name currState (storing the set of currently active states),
- the multi-set names $\text{inQueue}(O)$, $\text{outQueue}(O)$ (the input resp. output queue),
- the function name $\text{trigsusy}()$ mapping each operation name to the object or subsystem that last sent it (to allow sending back return values),
- the function name finished (mapping states to Boolean values, indicating whether a given state is finished), and
- all variables names in $\text{evt}(t)$ for all $t \in \text{Trs}_D$.

The Boolean finished_S is set to *true* at the end of an internal activity $\text{ActvRule}(\text{int}(S))$.

The ASM $\llbracket D \rrbracket^{SC}$ has two rules, **Initialize**(D) and **Main**(D), given below (defined using other rules omitted for space restrictions; they are adapted from (Börger et al., 2000) and can be found in (Jürjens, 2002)). The former rule initializes the variables of the ASM. The latter rule consists of selecting the event to be executed next (where priority is given to the completion event) and executing it, and then executing the rules for the internal activities in a random order.

Rule Initialize(D)

do – in – parallel

$\text{inQueue}(\text{Obj}_D) := \emptyset$ $\text{outQueue}(\text{Obj}_D) := \emptyset$

$\text{currState} := \{\text{Ini}_D\}$ $\text{finished}_{\text{Ini}_D} := \text{false}$

enddo

Rule Main(D)

seq if $\text{Completed} \neq \emptyset$ then $\text{eventExecution}(\text{CompEv})$

else choose e with $e \in \text{inQueue}(O)$ do

```

    seq inQueue( $O$ ) := inQueue( $O$ )  $\setminus$   $\{\{e\}\}$ 
      if  $e = op_{sender}[args] \in \text{Operation}$ 
        then seq  $e := op[args]$ 
          trigsusy( $e$ ) :=  $sender$  endseq
        eventExecution( $e$ )
      endseq
    loop  $S$  through set currState
      seq finished $_S$  := false
        ActvRule(int( $S$ ))
      endseq
    endseq
endseq

```

Here $\text{eventExecution}(e)$ is a macro whose execution models the execution of the event e , which involves firing its associated action as follows.

Rule **ActionRule**(call($op[args]$))
 $\text{outQueue}(O) := \text{outQueue}(O) \uplus \{\{op_O[args]\}\}$

Rule **ActionRule**(send(e))
 $\text{outQueue}(O) := \text{outQueue}(O) \uplus \{\{e\}\}$

Rule **ActionRule**(send(return(op)(a)))
 $\text{outQueue}(O) := \text{outQueue}(O) \uplus \{\{\text{trigsusy}(op).\text{return}(op)(a)\}\}$

3.2. Activity diagrams

An activity diagram is a special case of a state machine that is used to model processes involving one or more classes (UML, 2001, B-2). The internal activities in an activity diagram are given as an object or subsystem name S . Thus activity diagrams coordinate the execution of objects and subsystems. Note that these can be specified to proceed at different speeds, for example by including delay actions.

3.3. Subsystems and System Specifications

A subsystem $\mathcal{S} = (\text{name}(\mathcal{S}), \text{Msgs}(\mathcal{S}), \text{Ints}(\mathcal{S}), \text{Ssd}(\mathcal{S}), \text{Ad}(\mathcal{S}), \text{Bd}(\mathcal{S}))$ is given by

- the name $\text{name}(\mathcal{S})$ of the system,
- a set of accepted messages $\text{Msgs}(\mathcal{S})$,
- a set of interfaces $\text{Ints}(\mathcal{S})$,
- a static structure diagram $\text{Ssd}(\mathcal{S})$ (defined in Section 3.3)
- an activity diagram $\text{Ad}(\mathcal{S})$, and
- a set $\text{Bd}(\mathcal{S})$ of statechart diagrams (each for a different object). Each statechart specifies the behaviour of one object in the subsystem, so $\text{Bd}(\mathcal{S})$ is empty if there are no classes (but only subsystems) in $\text{Ssd}(\mathcal{S})$.

Subsystems may contain further kinds of diagrams omitted here.

For simplicity, we do not explicitly model creation and deletion of objects, but only activation and inactivation (following (Kleppe and Warmer, 2001, p. 15)). The run-to-completion step for each object is performed in parallel, subject to the flow of control specified by the activity diagram.

Class Diagrams or Static structure Diagrams. A class model $C = (cname, att, mess, int)$ is given by

- a name $cname$,
- a set of attribute names att ,
- a set of message names $mess$,
- and a set int of interfaces of the form $I = (iname, msg)$ where $iname$ is the interface name and msg a set of message names.

A *dependency* is a tuple $(dep, indep, int, stereo)$ consisting of

- class or subsystem names dep and $indep$,
- an interface name int , and
- a stereotype $stereo \in \{\langle\langle call \rangle\rangle, \langle\langle send \rangle\rangle\}$.

Note that the existence of a dependency $(dep, indep, int, stereo)$ with $stereo \in \{\langle\langle call \rangle\rangle, \langle\langle send \rangle\rangle\}$ implies that the object dep knows of the object $indep$ (otherwise dep could not send a signal to or call the object $indep$).

A *static structure diagram* $D = (SuSys(D), Dep(D))$ is given by a set $SuSys(D)$ consisting of class models or subsystems, and a set $Dep(D)$ of dependencies $(dep, indep, int, stereo)$. Other modelling elements (such as associations) can be added without complication. A static structure diagram is called a *class diagram* if it contains no subsystems.

Consistency between UML diagrams. A subsystem is called *consistent* if the following conditions are satisfied:

- The object and subsystem names appearing as activities in the activity diagram are part of the static structure diagram. The behaviour of the objects is defined by statecharts.
- For each call resp. send action in a statechart diagram, the static structure diagram C must have a dependency stereotyped $\langle\langle call \rangle\rangle$ resp. $\langle\langle send \rangle\rangle$ between the objects (or their interfaces) in question. For each assignment action $att := exp$ in S , att is contained in the set of attributes of Cls_S given in D .
- The operations offered by the subsystem must be offered by class models or subsystems in the static structure diagram.
- For statecharts S and T we have $S = T$ or $Obj_S \neq Obj_T$.

Behavioural semantics of Subsystems. Suppose we are given a consistent subsystem S . The behavioural interpretation of S is given by the ASM $\llbracket S \rrbracket^{SuSy}$ with the rules **Initialize**(S) and **Main**(S) defined in the following.

We write $\text{Acts}(S)$ for the set of object and subsystem names giving the activities in $\text{Ad}(S)$. If $act \in \text{Acts}(S)$ is an object whose behaviour is given by a statechart diagram $D \in \text{Bd}(S)$ (i. e. $\text{Obj}_D = act$), we write $\llbracket S \rrbracket_{act}^{BD}$ for $\llbracket D \rrbracket^{SC}$, $\text{BDInitialize}_S(act)$ for **Initialize**(D) and $\text{BDMain}_S(act)$ for **Main**(D). If $act \in \text{Acts}(S)$ is a subsystem in $\text{Ssd}(S)$, we write $\llbracket S \rrbracket_{act}^{BD}$ for $\llbracket act \rrbracket^{SuSy}$, $\text{BDInitialize}_S(act)$ for **Initialize**(act) and $\text{BDMain}_S(act)$ for **Main**(act).

The vocabulary of $\llbracket S \rrbracket^{SuSy}$ includes the following names:

- the names in the vocabularys of $\llbracket S \rrbracket^{SuSy}$ for all $S \in \text{SuSys}(\text{Ssd}(S))$, $\llbracket \text{Ad}(S) \rrbracket^{AD}$ and $\llbracket \text{Bd}(S) \rrbracket_S^{BD}$ for all $S \in \text{Acts}(S)$,
- the names $\text{inQueue}(S)$ and $\text{outQueue}(S)$ (in and out queue of S).

We define the two rules.

Rule Initialize(S)

```
do – in – parallel
  Initialize(Ad(S))
  forall S with S ∈ Acts(S) do
    BDInitializeS(S)
    inQueue(S) := ∅    outQueue(S) := ∅
  enddo
```

Rule Main(S)

```
seq
  forall S with S ∈ Acts(S) do
    inQueue(S) := inQueue(S) ⊔
      { tail(e) : e ∈ (inQueue(S) \ Msgs(S)) ∧ head(e) = S }
    inQueue(S) := ∅
    Main(Ad(S))
    forall S with S ∈ Acts(S) do
      inQueue(S) := inQueue(S) ⊔
        { tail(e) : e ∈ outQueue(T) ∧ head(e) = S }
      outQueue(S) := outQueue(S) ⊔
        { tail(e) : e ∈ outQueue(T) ∧ head(e) = S }
    forall S with S ∈ Acts(S) do
      outQueue(S) := ∅
    endseq
  endseq
```

System specifications. A UML system is given by a UML subsystem which is intended to model the complete system under consideration (rather than just a part).

Given a system \mathcal{S} and a sequence $\vec{I}_1, \dots, \vec{I}_n$ of multi-sets, the timed behaviour $\llbracket \mathcal{S} \rrbracket^t(\vec{I}_1, \dots, \vec{I}_n)$ is defined to be the set of possible contents of $\text{outlist}(\mathcal{S})$ after the execution of any instantiation of the following ASM rule on the ASM $\llbracket \mathcal{S} \rrbracket^{SuSy}$.

```

Rule  $\mathbf{tSys_A}(\mathcal{S})$ 
seq outlist( $\mathcal{S}$ ) :=  $\emptyset$ 
  Initialize( $\mathcal{S}$ )
  loop  $i$  through list  $[1 \dots n]$ 
    seq inQueue( $\mathcal{S}$ ) := inQueue( $\mathcal{S}$ )  $\uplus$   $\vec{I}_i$ 
      Main( $\mathcal{S}$ )
      outlist( $\mathcal{S}$ ) := outlist( $\mathcal{S}$ ).outQueue( $\mathcal{S}$ )
      outQueue( $\mathcal{S}$ ) :=  $\emptyset$ 
    endseq
  endseq
endseq

```

The untimed behaviour of a system \mathcal{S} is a function $\llbracket \mathcal{S} \rrbracket^u()$ from multi-sets of events (the inputs to the system during its execution) to sets of multi-sets of events (the possible outputs of the system during its (possibly non-terminating) execution). It is defined at the meta-level. Given a multi-set I , we define $\llbracket \mathcal{S} \rrbracket^u(I)$ to be the fix-point of the content of $\text{outQueue}(\mathcal{S})$ after firing the rule

```
seq Initialize( $\mathcal{S}$ ) inQueue( $\mathcal{S}$ ) :=  $I$  endseq
```

and iterating the rule **Main**(\mathcal{S}). The fix-point exists, because \mathcal{S} changes $\text{outQueue}(\mathcal{S})$ only by *adding* elements: it is the directed union of the contents of $\text{outQueue}(\mathcal{S})$ after each execution of **Main**(\mathcal{S}).

4. Equivalence and Refinement

Definition 1 (Property Refinement) Suppose we are given two subsystem specifications \mathcal{S} and \mathcal{S}' and a set \mathcal{M} of message names. \mathcal{S}' is a timed (resp. untimed) \mathcal{M} -refinement of \mathcal{S} if conditions I and IIa (resp. conditions I and IIb) are fulfilled:

I $\text{Msgs}(\mathcal{S}) \cap \mathcal{M} \subseteq \text{Msgs}(\mathcal{S}') \cap \mathcal{M}$

IIa for all sequences I_1, \dots, I_n of multi-sets of events such that $\text{msgname}(e) \in \mathcal{M}$ for each $e \in \bigcup_{i=1, \dots, n} I_i$, we have

$$\llbracket \mathcal{S}' \rrbracket^t(I_1, \dots, I_n) \curvearrowright \mathcal{M} \subseteq \llbracket \mathcal{S} \rrbracket^t(I_1, \dots, I_n) \curvearrowright \mathcal{M}$$

IIb for each multi-set I of events such that $\text{msgname}(e) \in \mathcal{M}$ for each $e \in I$, we have $\llbracket \mathcal{S}' \rrbracket^u(I) \curvearrowright \mathcal{M} \subseteq \llbracket \mathcal{S} \rrbracket^u(I) \curvearrowright \mathcal{M}$

where for a set S of multi-sets of events we define $S \curvearrowright \mathcal{M} \stackrel{\text{def}}{=} \{M \setminus \{e \in \mathbf{Events} : \text{msgname}(e) \in \mathcal{M}\} : M \in S\}$ to be the set of multi-sets obtained by removing all events e with $\text{msgname}(e) \notin \mathcal{M}$ (and extend the definition to sequences of multi-sets by applying it to the elements of the sequence). A *timed (resp. untimed) MsgNm-refinement* is simply called a *(resp. untimed) refinement*.

For a motivation of the refinements and comparison to other kinds, see (Jürjens, 2002). Both kinds of \mathcal{M} -refinement are reflexive and transitive for each \mathcal{M} . They preserve all \mathcal{M} -safety properties (i. e. sets of sequences of event multi-sets with names in \mathcal{M}). We show that refinement is preserved by substitution (i. e. a precongruence wrt. composition by subsystem formation).

A *parameterized subsystem* $\mathcal{S}(\mathcal{Y}_1, \dots, \mathcal{Y}_n)$ is a subsystem specification where n of the subsystems are replaced by subsystem variables \mathcal{Y}_i . For subsystems $\mathcal{S}_1, \dots, \mathcal{S}_n$, $\mathcal{S}(\mathcal{S}_1, \dots, \mathcal{S}_n)$ is the subsystem obtained by substituting \mathcal{S}_i for \mathcal{Y}_i , for each i , in \mathcal{S} .

Theorem 1 *If \mathcal{S}'_i is a timed refinement of \mathcal{S}_i for each $i = 1, \dots, n$ then for any parameterized subsystem $\mathcal{S}(\mathcal{Y}_1, \dots, \mathcal{Y}_n)$, $\mathcal{S}(\mathcal{S}'_1, \dots, \mathcal{S}'_n)$ is a timed refinement of $\mathcal{S}(\mathcal{S}_1, \dots, \mathcal{S}_n)$.*

The theorem does not hold for untimed refinements (Jürjens, 2002).

Definition 2 (Behavioural Equivalence) *Two subsystem specifications \mathcal{S} and \mathcal{S}' are timed (resp. untimed) behaviourally equivalent if $\text{Msgs}(\mathcal{S}) = \text{Msgs}(\mathcal{S}')$, \mathcal{S} is a timed (resp. untimed) $\text{Msgs}(\mathcal{S})$ -refinement of \mathcal{S}' , and \mathcal{S}' is a timed (resp. untimed) $\text{Msgs}(\mathcal{S})$ -refinement of \mathcal{S} .*

The above facts on timed refinement imply that timed behavioural equivalence is a congruence wrt. composition by subsystem formation.

Behavioural equivalence can be used e.g. to verify consistency of two subsystem specifications that are supposed to describe the same behaviour.

5. Example: Secure channel establishment

Figure 1 gives a high level system specification \mathcal{S} for communication from a sender object to a receiver object, including a class diagram with appropriate interfaces (we leave out the activity diagram which simply specifies that the two objects are executed in parallel).

Assume that an assessment of the physical layer of the system shows that security requirements are not provided. Thus we construct a refinement \mathcal{S}' in Figure 2. The class diagram needs to allow the receiver to send his certificate to the sender. The behaviour of the sender includes retrieving the public key and the certificate from the sender, checking the certificate, and encrypting the data.

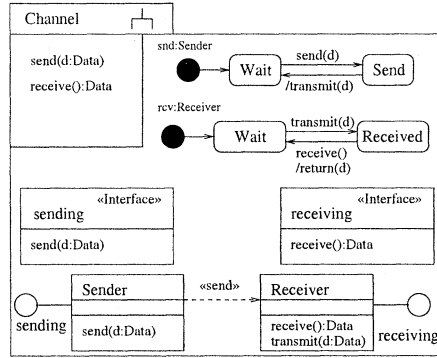


Figure 1. Sender and receiver

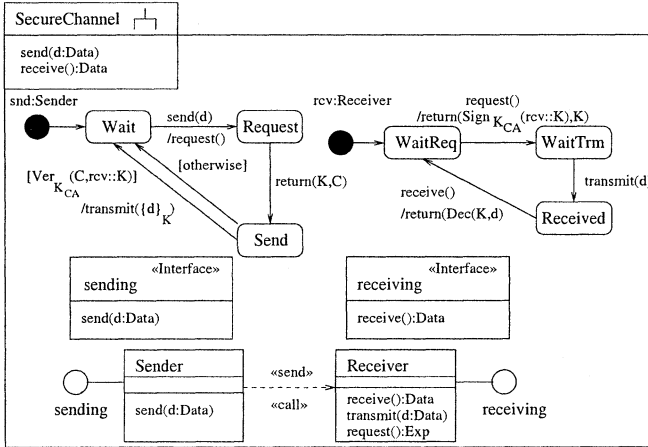


Figure 2. Secure channel

The receiver gives out the key and certificate first, and decrypts the received data. We assume that $\text{Ver}_K(\text{Sign}_K(m), m) = \text{true}$ (i. e. $\text{Ver}_K(m, O)$ is true if m is a valid signature of O with key K).

Theorem 2 \mathcal{S}' is an untimed $\{\text{send}, \text{receive}, \text{return}_{\text{receive}}\}$ -refinement of \mathcal{S} .

The (straightforward) proof can be found in (Jürjens, 2002). Note that the theorem does not imply that the refined protocol is secure; this has to be established by other means (including an extension of the formal semantics with adversary scenarios), see (Jürjens, 2002).

Note that \mathcal{S}' is not a *timed* $\{\text{send}, \text{receive}, \text{return}_{\text{receive}}\}$ -refinement of \mathcal{S} (because of the delay caused by the key exchange).

6. Related Work

There has been a considerable amount of work towards a formal semantics for various parts of UML (see (Jürjens, 2002) for an overview). (Evans et al., 1999) discusses some fundamental issues concerning a formal foundation for UML. (Börger et al., 2000) uses ASMs to give a semantics for statecharts. (Stevens, 2001) gives a semantics for use case diagrams based on the process algebra CCS. (Övergaard and Palmkvist, 2000) considers interacting UML subsystems, but without giving a formal semantics. A combined formal semantics for UML statecharts and class diagrams has been given in (Reggio et al., 2001). (Schäfer et al., 2001) gives a semantics for statecharts and shows exemplarily how to check whether a set of statecharts satisfies a collaboration. Refinements have been investigated in the object-oriented setting e.g. in (Derrick and Smith, 2000; Derrick and Boiten, 2001) where the introduced structural refinement has a similar motivation as our interface refinement. In the context of subtyping, refinement has been considered e.g. in (Poetzsch-Heffter, 1997).

7. Conclusion and Future Work

Our work shows that giving a formal foundation for complete UML specifications, rather than just single diagrams, is possible. While in our presentation here we left out some of the more advanced features of some of the considered diagrams, incorporating them is not a problem (except for an increase in complexity). This paves the way for unambiguous modelling with UML, and in particular for tool-support that can simulate specifications in their entirety. This would seem to be very useful indeed, since one of the main challenges in constructing large-scale software systems is to ensure that different components act together as expected.

ASMs proved to be a quite adequate tool to handle the complexities in formally combining the different UML diagrams, due to their flexibility. In particular, the crucial extensions over the semantics for statecharts in (Börger et al., 2000) could be done in a rather natural way.

The notions of subsystem equivalence and refinement seem to be valuable tools in the construction and analysis of UML specification (although here we could only present a toy example).

On the application side, the semantics presented here has been useful in the context of security-critical systems (cf. e.g. (Jürjens, 2001-b)).

In further work (Jürjens, 2002) we have treated other kinds of UML diagrams and other kinds of refinements (including one inspired by action refinement (Gorrieri and Rensink, 2000)). We aim to give conditions for checking equivalence and refinement on the level of syntax or by combining an object-based logic (e.g. (Distefano et al., 2000)) with a logic for ASMs (e.g. (Poetzsch-Heffter, 1994)).

In terms of tool-support we intend to extend the ASM-based statechart simulator presented in (Campo et al., 2001) to the semantics presented here.

Acknowledgements. Discussions with A. Cavarra about formal semantics for UML and constructive suggestions by the anonymous referees to improve the presentation of the paper are gratefully acknowledged.

References

- Börger, E., Cavarra, A. and Riccobene, E. (2000). Modeling the dynamics of UML State Machines. In Y. Gurevich, P. Kutter, M. Odersky, and L. Thiele, editors, *Abstract State Machines. Theory and Applications*, volume 1912 of *LNCS*, pages 223–241. Springer-Verlag.
- Broy, M. and Stølen, K. (2001). *Specification and Development of Interactive Systems*. Springer-Verlag.
- Campo, R., Cavarra, A. and Riccobene, E. (2001). Simulating UML state machines. In E. Börger and U. Glässer, editors, *ASM'2001*, *LNCS*. Springer-Verlag. To be published.
- Derrick, J. and Boiten, E. (2001). *Refinement in Z and Object-Z: Foundations and advanced applications*. Formal Approaches to Computing and Information Technology. Springer-Verlag.
- Distefano, D., Katoen, J.-P. and Rensink, A. (2000). On a temporal logic for object-based systems. In Smith and Talcott (Smith and Talcott, 2000), pages 305–326.
- Derrick, J. and Smith, G. (2000). Structural refinement in Object-Z / CSP. In W. Grieskamp, T. Stanten, and B. Stoddart, editors, *Integrated Formal Methods (IFM 2000)*, volume 1945 of *LNCS*, pages 194–213. Springer-Verlag.
- Evans, A., France, R., Lano, K. and Rumpe, B. (1999). The UML as a formal modeling notation. In J. Bézivin and P.-A. Muller, editors, *The Unified Modeling Language - Workshop UML'98: Beyond the Notation*, *LNCS*, pages 297–307. Springer-Verlag.
- Gorrieri, R. and Rensink, A. (2000). Action refinement. In J. Bergstra, A. Ponse, and S. Smolka, editors, *Handbook of Process Algebra*. Elsevier.
- Gurevich, Y. (1995). Evolving algebras 1993: Lipari guide. In E. Börger, editor, *Specification and Validation Methods*, pages 9–36. OUP.
- Hußmann, H., editor. (2001). *Fundamental Approaches to Software Engineering (FASE, 4th International Conference)*, volume 2029 of *LNCS*. Springer-Verlag.
- Jürjens, J. (2001-a). Secrecy-preserving refinement. In *Formal Methods Europe (International Symposium)*, volume 2021 of *LNCS*, pages 135–152. Springer-Verlag.
- Jürjens, J. (2001-b). Towards development of secure systems using UMLsec. In Hußmann (Hußmann, 2001), pages 187–200.
- Jürjens, J. (2002). *Principles for Secure Systems Design*. PhD thesis, Oxford University Computing Laboratory. Submitted.
- Kleppe, A. and Warmer, J. (2001). Unification of Static and Dynamic Semantics of UML.
- Övergaard, G. and Palmkvist, K. (2000). Interacting Subsystems in UML. In A. Evans, S. Kent, and B. Selic, editors, *The Unified Modeling Language: Advancing the Standard (UML'2000)*, volume 1939 of *LNCS*. Springer-Verlag.
- Poetzsch-Heffter, A. (1994). Deriving partial correctness logics from evolving algebras. In B. Pehrson and I. Simon, editors, *IFIP 13th World Computer Congress '94*, pages 434–439. Elsevier.
- Poetzsch-Heffter, A. (1997). Specification and verification of object-oriented programs. Habilitation thesis, Technical University of Munich.

- Reggio, G., Cerioli, M. and Astesiano, E. (2001). Towards a Rigorous Semantics of UML Supporting its Multiview Approach . In Hußmann (Hußmann, 2001).
- Rumbaugh, J., Jacobson, I. and Booch, G. (1999). *The Unified Modeling Language Reference Manual*. Addison-Wesley.
- Schäfer, T., Knapp, A. and Merz, S. (2001). Model checking UML state machines and collaborations. In S.D. Stoller and W. Visser, editors, *Workshop on Software Model Checking*, volume 55 of *ENTCS*. Elsevier.
- Stevens, P. and Pooley, R. (2000). *Using UML*. Addison-Wesley.
- Stärk, R., Schmid, J. and Börger, E. (2001). *Java and the Java Virtual Machine – Definition, Verification, Validation*. Springer-Verlag.
- Smith, S. F. and Talcott, C. L., editors. (2000). *4th International Conference on Formal Methods for Open Object-Based Distributed Systems (FMOODS 2000)*. IFIP TC6/WG6.1, Kluwer Academic Publishers.
- Stevens, P. (2001). On use cases and their relationships in the Unified Modelling Language. In Hußmann (Hußmann, 2001), pages 140–155.
- UML Revision Task Force. (2001). OMG UML Specification v. 1.4. OMG Document ad/01-09-67. Available at [http : //www.omg.org/uml](http://www.omg.org/uml).