

# An Automata-Theoretic Framework for Real-Time Actors Schedulability and Compatibility

MOHAMMAD MAHDI JAGHOORI, CWI, Amsterdam, The Netherlands

DELPHINE LONGUET, University Paris-Sud, LRI UMR8623, Orsay, F-91405, Paris, France

FRANK DE BOER, CWI, Amsterdam, The Netherlands

TOM CHOTHIA, School of Computer Science, University of Birmingham, UK

MARJAN SIRJANI, Reykjavík University, Iceland; University of Tehran and IPM, Tehran, Iran

Schedulability analysis in real time systems amounts to checking whether every task is processed within its designated deadline. We introduce a modular automata-theoretic approach to the schedulability analysis of actor-based distributed systems in order to avoid the state-space explosion problem. Behavioral interfaces are key to modularity by specifying deadlines and timings for messages: every actor (seen as an off-the-shelf module) is to guarantee the deadlines on its tasks (triggered by the incoming messages), provided the messages arrive as expected. To be schedulable, it is essential that actors can apply customized scheduling policies. We then propose a novel method to test the compatibility of individually schedulable actors (with respect to their behavioral interfaces) in terms of a refinement notion extended to take deadlines into account. We show that the analyses are decidable and automate the process using the UPPAAL model checker.

Categories and Subject Descriptors: D.2.4 [Software Engineering]: Software Verification—*formal methods; model checking*; D.2.5 [Software Engineering]: Testing and Debugging—*testing strategies*; D.4.1 [Operating Systems]: Process Management—*scheduling*; D.4.7 [Operating Systems]: Organization and Design—*Real-time systems and embedded systems*

General Terms: Languages, Theory, Verification

Additional Key Words and Phrases: Actor model, Concurrent objects, Real-time, Scheduling, Timed automata

## 1. INTRODUCTION

Actors were originally introduced by Hewitt as autonomous reasoning objects [Hewitt 1971]. Actor languages have then evolved as a powerful computational tool for modeling distributed and concurrent systems [Agha 1990; Agha et al. 1997]. Different extensions of Actors are proposed in several domains and are claimed to be the suitable model of computation for the most dominating applications [Hewitt 2007]. Examples of these domains include designing embedded systems [Lee et al. 2003; Lee et al. 2009], wireless sensor networks [Cheong et al. 2005], multi-core programming [Karmani et al. 2009] and designing web services [Chang and Agha 2007b; 2007a].

In an Actor-based model, actors are (re)active objects with encapsulated data and methods which represent their state and behavior, respectively. Actors are the units of concurrency, i.e., an actor conceptually has a dedicated processor. In the pure asynchronous setting [Hewitt 1971; Agha 1990], actors can only send asynchronous messages and have queues for

---

This work was partly funded by the European IST-33826 STREP project CREDO and FP7-231620 project HATS. The work of the second author was carried out during the tenure of an ERCIM “Alain Bensoussan” Fellowship Program.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© YYYY ACM 0164-0925/YYYY/01-ARTA \$10.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

receiving messages. An actor progresses by taking a message out of its queue and processing it by executing its corresponding method. A method is a piece of sequential code that may send messages. We may use the terms actors and objects interchangeably.

This model of concurrent computation forms the basis of the programming languages Erlang [Armstrong 2010] and Scala [Haller and Odersky 2009] that have recently gained in popularity, in part due to their support for scalable concurrency. However, for optimal use of both hardware and software resources, we cannot avoid leveraging scheduling and performance related issues from the underlying operating system to the application level as argued for example in [Berman and Wolski 1996; Alur and Weiss 2008]. In general, the goal of *resource-aware programming* (RAP [Moreau and Queinnec 2005], SUMATRA [Acharya et al. 1996], CAMELOT [MacKenzie and Wolverson 2003]) is to express policies for the management of resources. In this paper, we focus on CPU time and extend the actor model with real-time, e.g., deadlines associated to messages, and explicit application-level scheduling policies associated to the individual actors (rather than, for instance, assuming “First Come First Served” (FCFS) by default).

In this paper we introduce a general framework for schedulability analysis of abstract Actor models enhanced with real-time and scheduling policies (see Fig. 1). The real-time specification of an actor, consisting of its methods, scheduler and queue, is given in timed automata [Alur and Dill 1994]. A scheduling policy determines the order in which the (methods corresponding to) queued messages should be executed. We restrict to schedulers in which a new message cannot preempt the currently running method. Method automata mainly specify what messages are sent and when. Receiving messages are handled by the schedulers and can be seen as events generating tasks. A deadline is assigned to each message specifying the time before which the intended job should be accomplished. This framework allows quality-of-service and deployment requirements to be analyzed and resolved at design time. For example in a real-time setting, we must guarantee a maximum on average response-time (end-to-end deadlines) or a minimum on the level of system throughput. For our analysis, we have chosen to use UPPAAL [Larsen et al. 1997]; this choice is however not essential and in principle other tools for timed automata can also be employed. This framework complements the work in [Nobakht et al. 2012], in which we describe how application-level scheduling policies can be implemented into a programming language on top of Java.

As seen in Fig. 1, a behavioral interface provides an abstract and high-level view of the actor by abstracting from the queue and method implementations. In fact, it shows the pattern of possible interactions it may have with its environment. This captures the valid sequences of the provided and required services (received and sent messages). We model a behavioral interface as a deterministic timed automaton to further capture the timings and deadlines of the messages. For instance, the behavioral interface of a server that handles at most one request at a time can be defined as a loop of receiving a ‘request’ followed by a ‘reply’, i.e., no second request is allowed before providing the reply.

Fig. 1 further illustrates the different concepts underlying our modular schedulability analysis technique. In previous work [Jaghoori et al. 2009], based on the ideas of Task Automata [Fersman et al. 2007], we have shown how to analyze an individual actor with respect to its behavioral interface. As an actor is expected to be used as specified by its behavioral interface, we restrict the actor behavior accordingly when checking its schedulability in isolation. Our method allows us to statically find an upper bound on the length of schedulable queues; hence, the behavioral model of the actor has a finite number of states and the analysis is decidable. This way, one can analyze the actor with regard to different scheduling strategies, and find the best strategy.

The main contribution of this paper is the generalization of the schedulability analysis of individual actors to distributed systems of interacting actors. Once an actor is proved schedulable, in order to use it as an off-the-shelf component we need to check additionally that its actual usage in a given distributed system follows the expected usage (as speci-

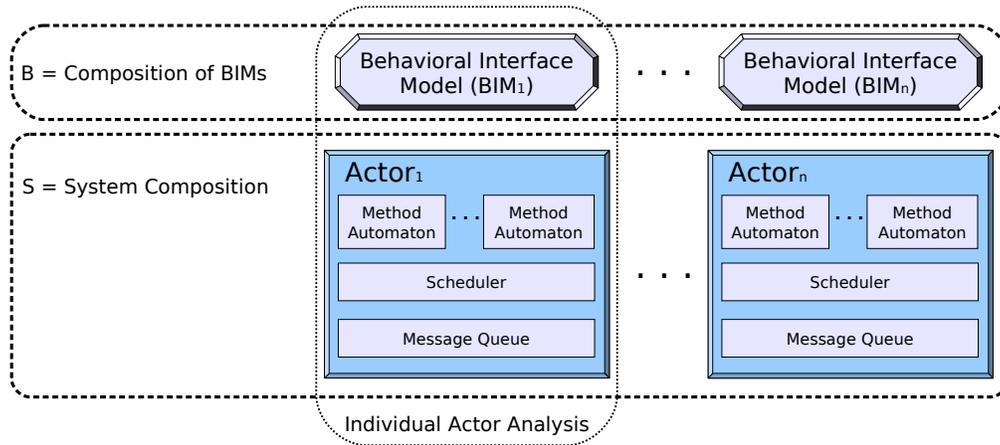


Fig. 1. An off-the-shelf actor component is guaranteed to be schedulable if it is used as expected in its behavioral interface. This correct usage in a system  $S$  can be tested, called the compatibility check.

fied by the behavioral interface), called the compatibility check. In other words, checking compatibility for each actor  $A_i$  involves ensuring that the rest of the system respects the requirements specified in its behavioral interface  $B_i$ . Unfortunately, as we explain in Section 4, this cannot be checked simply at the level of behavioral interfaces, because they include only an approximation of when messages are sent.

Theoretically, one approach to checking compatibility is to construct the complete system behavior  $S$  of all actors together and show that  $S$  is a *refinement* of the product of the automata for behavioral interfaces (call it  $B$ ), i.e., the set of timed traces of  $S$  is a subset of that of  $B$ . Assuming that  $B$  is deterministic, in theory one can prove that  $S$  is refinement of  $B$  by model-checking the synchronous product of  $S$  and  $B$  (restricted by the computed queue bounds). Due to the message queues (one for each actor) in  $S$ , this would lead however to an unmanageable state-space explosion. Instead of verifying the refinement relation, we introduce a novel method for *counter-example oriented* testing. In this method we generate a test case from  $B$  as follows. We take a trace from  $B$  and complete it into a test case for  $S$  by adding transitions that capture all possible one-step deviations from the original trace. Among these transitions, those not allowed in  $B$  produce a counter-example, i.e., a *trace of  $S$  which does not belong to  $B$* . This technique is much more effective than generating test cases from  $S$  to be checked against  $B$ , because it allows for automated generation of test cases from  $B$  (note that  $B$  does not involve queues) and a reduction of the overall system behavior  $S$  by the test case.

**Contribution.** As mentioned above, the overall contribution of this paper is the integration of our previous work, i.e., schedulability analysis of individual actors [Jaghoori et al. 2009] and a preliminary report on testing compatibility [Jaghoori et al. 2008], into a complete framework for modular schedulability analysis of real-time systems of actors. This paper further provides a more precise notion of refinement including quiescence and deadlines as well as a formal justification of the testing technique described above. More specifically, we provide the proofs for soundness and completeness. Moreover, we introduce in this paper the notion of rigidity: A test case is *rigid* if and only if it detects any incorrect refinement along the traces of the test case. We also prove rigidity. We illustrate the applicability of this framework by modeling and analyzing a hybrid peer-to-peer architecture, i.e., peer nodes are connected by a central server, called the broker. The network configuration changes dynamically as the broker connects different peers. It is worthwhile to observe

that the testing approach is still applicable in presence of dynamic reconfiguration. This is because the automata nature of behavioral interfaces can capture this dynamic behavior.

**Paper Structure.** Section 2 provides the grounds for the approach by explaining timed automata. In Section 3, we explain how we model real-time actors and exemplify this by means of a peer-to-peer case study. A review of our modular schedulability analysis technique is given in Section 4. Section 5 describes our approach to testing refinement for timed automata, which is then applied to test compatibility in the context of schedulability analysis. Related works are presented in Section 6. Section 7 concludes the paper.

## 2. PRELIMINARIES: TIMED AUTOMATA

We base our techniques on timed automata and thus can take advantage of the abundant tools available. As we choose UPPAAL, we tailor the definitions accordingly.

*Syntax.* Let  $Act$  be a finite set of *actions*. Let  $C$  be a finite set of real-valued *clocks*. We define  $\mathcal{B}(C)$ , the set of clock constraints, as the set of boolean formulas built over elementary constraints  $x \sim n$  and  $x - y \sim n$  where  $x, y \in C$ ,  $n \in \mathbb{N}$ , and  $\sim \in \{<, \leq, =, \geq, >\}$ , with boolean operators  $\vee$ ,  $\wedge$  and  $\neg$ . A *timed automaton*  $A$  over  $Act$  and  $C$  is a tuple  $(L, l_0, E, I)$  where  $L$  is a finite set of *locations*, with  $l_0 \in L$  being the initial location.  $E \subseteq L \times \mathcal{B}(C) \times Act \times 2^C \times L$  is a finite set of *edges*. We write  $l \xrightarrow{g, a, r} l'$  for an edge from location  $l$  to location  $l'$  guarded by a clock constraint  $g \in \mathcal{B}(C)$ , labeled with the action  $a \in Act$  and resetting the subset  $r$  of  $C$ . Finally,  $I : L \rightarrow \mathcal{B}(C)$  assigns an *invariant* to each location. Location invariants are restricted to conjunctions of constraints of the form  $x < n$  or  $x \leq n$  for  $x \in C$  and  $n \in \mathbb{N}$ .

*Semantics.* A timed automaton defines an infinite labeled transition system whose states are pairs  $(l, u)$  where  $l \in L$  and  $u : C \rightarrow \mathbb{R}_+$  is a *clock assignment*. We denote by  $\mathbf{0}$  the assignment mapping every clock in  $C$  to 0. The initial state is  $s_0 = (l_0, \mathbf{0})$ . There are two types of transitions: action transitions  $(l, u) \xrightarrow{a} (l', u')$  where  $a \in Act$ , if there exists  $l \xrightarrow{g, a, r} l'$  such that  $u$  satisfies the guard  $g$ ,  $u'$  is obtained by resetting to zero all clocks in  $r$  and leaving the others unchanged and  $u'$  satisfies the invariant of location  $l'$ ; delay transitions  $(l, u) \xrightarrow{d} (l, u')$  where  $d \in \mathbb{R}_+$ , if  $u'$  is obtained by delaying every clock for  $d$  time units and for each  $0 \leq d' \leq d$ ,  $u'$  satisfies the invariant of location  $l$ .

For a sequence of labels  $w = w_1 w_2 \dots w_n$ , we write  $s_0 \xrightarrow{w} s_n$  to denote the sequence of transitions  $s_0 \xrightarrow{w_1} s_1 \rightarrow \dots \xrightarrow{w_n} s_n$ .

*Deterministic timed automata.* A timed automaton is called *deterministic* if and only if given any two edges  $l \xrightarrow{g, a, r} l'$  and  $l \xrightarrow{g', a, r'} l''$ , the guards  $g$  and  $g'$  are disjoint (i.e.  $g \wedge g'$  is unsatisfiable). Furthermore, there is at most one transition with an invisible action  $l \xrightarrow{g, \tau, r} l'$  from any location  $l$ , in which case,  $g$  is disjoint from guards of other transitions from  $l$ .

*Variables.* As accepted in UPPAAL, we allow variables of type boolean and bounded integers for each automaton. Variables can appear in guards and updates. The semantics of timed automata changes such that each state will include the current values of the variables as well, i.e.  $(l, u, v)$  with  $v$  a variable assignment. An action transition  $(l, u, v) \xrightarrow{a} (l', u', v')$  additionally requires  $v$  and  $v'$  to be considered in the corresponding guard and update.

*Timed automata with inputs and outputs.* In the following, we assume the set of actions  $Act$  is partitioned into two disjoint sets: a set  $Act_I$  of *input actions*  $a?$  and a set  $Act_O$  of *output actions*  $a!$ . A *non-observable internal action*  $\tau$  is also assumed. A timed automaton with inputs and outputs is a timed automaton over  $Act_\tau = Act \cup \{\tau\}$ .

*Networks of timed automata.* A system may be described as a collection of timed automata with inputs and outputs  $A_i$  ( $1 \leq i \leq n$ ) communicating with each other. The behavior of

the system is then defined as the parallel composition of those automata  $A_1 \parallel \dots \parallel A_n$ . Semantically, the system can delay if all automata can delay and can perform an action if one of the automata can perform an internal action or if two automata can synchronize on complementary actions (inputs and outputs are complementary). In a network of timed automata, variables can be defined locally for one automaton, globally (shared between all automata), or as parameters to the automata.

A location can be marked *urgent* in an automaton to indicate that the automaton cannot spend any time in that location. This is equivalent to resetting a fresh clock  $x$  in all of its incoming edges and adding an invariant  $x \leq 0$  to the location. In a network of timed automata, the enabled transitions from an urgent location may be interleaved with the enabled transitions from other automata (while time is frozen). Like urgent locations, *committed* locations freeze time; furthermore, if any process is in a committed location, the next step must involve an edge from one of the committed locations.

*Timed traces.* A *timed sequence*  $\sigma \in (Act_\tau \cup \mathbb{R}_+)^*$  is a sequence of timed actions in the form of  $\sigma = t_1 a_1 t_2 a_2 \dots a_n t_{n+1}$  such that for all  $i$ ,  $1 \leq i \leq n$ ,  $t_i \leq t_{i+1}$ . Given a timed sequence  $\sigma$ ,  $\pi_{obs}(\sigma)$  denotes the projection of  $\sigma$  on  $Act$ , intuitively deleting  $t_i \tau$  occurrences. The sequence  $\pi_{obs}(\sigma)$  is called the *observable timed sequence* associated to  $\sigma$ .

A *run* of a timed automaton  $A$  from initial state  $(l_0, \mathbf{0})$  over a timed sequence  $\sigma = t_1 a_1 t_2 a_2 \dots a_n t_{n+1}$  is a sequence of transitions:

$$(l_0, \mathbf{0}) \xrightarrow{d_1} (l_0, u'_0) \xrightarrow{a_1} (l_1, u_1) \xrightarrow{d_2} \dots \xrightarrow{a_n} (l_n, u_n) \xrightarrow{d_{n+1}} (l_n, u'_n)$$

where  $d_1 = t_1$  and for all  $i$ ,  $1 < i \leq n + 1$ ,  $t_i = t_{i-1} + d_i$ . The set  $Traces(A)$  of timed traces of  $A$  is the set of timed sequences  $\sigma$  for which there exists a run of  $A$  over  $\sigma$ . The set  $Traces_{obs}(A)$  of observable timed traces of  $A$  is the set  $\{\pi_{obs}(\sigma) \mid \sigma \in Traces(A)\}$ .

### 3. ACTORS AS REAL-TIME ASYNCHRONOUS CONCURRENT OBJECTS

We describe in this section how to use automata theory, along the lines of our previous work [Jaghoori et al. 2009; Jaghoori et al. 2008], to describe actors. Actors in our framework specify local scheduling strategies, e.g., based on fixed priorities, earliest deadline first, or a combination of such policies. Real-time actors may need certain customized scheduling strategies in order to meet their QoS requirements. Our approach can be easily adapted to any actor-based modeling platform, e.g., Rebeca [Sirjani et al. 2004], Creol [Johnsen and Owe 2007], etc.

#### 3.1. A Formal Model of Actors

We present the semantics of an actor at two levels of abstraction (cf. Fig. 1). First a synthetic abstract behavior of the actor is given in one place, called its behavioral interface. Second, a more detailed specification of the actor behavior is given in terms of its methods, each modeled as an automaton. Finally, we describe how composition of multiple actor instances comprises a closed system.

*Behavioral Interface Model.* A behavioral interface specifies at a high level, and in the most general terms, how an actor behaves. It consists of the messages an actor may receive and send. A behavioral interface abstracts from specific method implementations, the message queue in the actor and the scheduling strategy. As explained later in this section, behavioral interfaces are key to modular analysis of actors.

To formally define a behavioral interface, we assume a finite set  $\mathcal{M}$  for method names. A behavioral interface  $B$  providing a set of method names  $M_B \subseteq \mathcal{M}$  is a deterministic timed automaton over alphabet  $Act^B$  such that  $Act^B$  is partitioned into two sets of actions:

— outputs:  $Act_O^B = \{m! \mid m \in \mathcal{M} \wedge m \notin M_B\}$

— inputs:  $Act_I^B = \{m(d)? \mid m \in M_B \wedge d \in \mathbb{N}\}$

The number  $d$  associated to input actions represents a deadline. A correct implementation of the actor should be able to finish method  $m$  before  $d$  time units. We are restricted to natural numbers for deadlines, because using real numbers makes analysis of timed automata undecidable. Output actions are the methods called by this actor and should be handled by (other actors in) the environment.

The semantics of a behavioral interface is defined simply as the timed traces on its action set. We define composition of behavioral interfaces as their synchronous product on complementary actions, where an output action  $m!$  synchronizes with input actions  $m(d)?$  and produces the action  $m(d)$  in the composed automaton.

Behavioral interfaces conceptually serve two purposes: 1) represent the actor to the environment, as explained above; 2) represent the environment to the actor. The latter function can be enabled by syntactically swapping the  $!$  and  $?$  signs in a behavioral interface. Thus one obtains an abstraction of the environments in which an actor may be used. In the following sections, this abstraction will be used for modeling and analysis of actors in UPPAAL.

*Actor Definition.* An actor may implement a behavioral interface  $B$  by providing implementation for the methods in  $M_B$ . An actor  $R$  implementing the behavioral interface  $B$  is a set  $\{(m_1, A_1), \dots, (m_n, A_n)\}$  where:

- $M_R = \{m_1, \dots, m_n\} \subseteq \mathcal{M}$  is a set of method names such that  $M_B \subseteq M_R$ ; and,
- for all  $i$ ,  $1 \leq i \leq n$ ,  $A_i$  is a timed automaton representing method  $m_i$  with the alphabet  $Act_i = \{m! \mid m \in M_R\} \cup \{m(d)! \mid m \in \mathcal{M} \wedge d \in \mathbb{N}\}$

There is at least a method `initial` in each actor, which is responsible for initialization. Method automata only send messages while computations are abstracted into time delays. Sending a message  $m \in M_R$  is called a self call. A self call with no explicit deadline inherits the (remaining) deadline of the method that triggers it; this mechanism is called delegation. Other send operations are to be given explicit deadlines.

Receiving messages and queuing them is done implicitly. Each actor has an unbounded queue for storing its incoming messages. For each message, a queue needs to store the method name  $m_i$  and its deadline  $d_i$ . Furthermore, it needs a clock  $c_i$  to keep track of its waiting time in the queue. This clock is reset to zero when the message is added to the queue. This message misses its deadline when  $c_i > d_i$ . Later we show how to put a reasonable bound on the queue size for analysis and therefore a finite number of clocks are needed. We assume that the first message in the queue represents the currently running method.

The semantics of an actor can be defined as a timed automaton, called *actor automaton*. Every location of the actor automaton is written as a pair  $(l, q)$  where  $q$  represents the contents of the queue and  $l$  refers to the current location of the currently executing method, i.e., the first method in the queue. The actor takes one transition if the currently running method takes a step. On the other hand, changes to the queue also cause a transition in the actor automaton. The latter type of transitions depend on the composition of the actor in the context of an environment, e.g. a closed system as defined below.

*System Composition.* In a system, a number of actors run concurrently, each maintaining a queue of the message it has to process. First of all, every actor instance must be associated with a specific scheduler. Messages are inserted into the queue by a scheduler in the order they should be executed, based on a scheduling strategy, e.g., FCFS (First Come First Served) or EDF (Earliest Deadline First). Typically the scheduler could dynamically examine the remaining time before the deadline of each message in the queue. However, to be able to statically write down the specification of a scheduler, we define a scheduler function that returns the set of all possibilities for putting the new message in the queue depending on different clock values. Formal definition of a scheduler is given in the ap-

A ‘first come first served’ scheduler always puts the new job at the back of the queue. Given a queue  $q = [m_1(d_1, c_1), \dots, m_k(d_k, c_k)]$  and a new task  $m(d)$ , we write this as:

$$\mathbf{FCFS}(\mathbf{q}, \mathbf{m}(d)) = \left\{ (true \implies [m_1(d_1, c_1), \dots, m_k(d_k, c_k), \underline{m(d, c)}]) \right\}$$

where  $c$  is a fresh clock not assigned to any message in  $q$ .

An ‘earliest deadline first’ scheduler inserts messages into the queue based on the remaining deadlines of the existing queue members. The scheduler cannot reorder the queue, so here we assume as an invariant that the messages already in the queue are in the right order.

$$\mathbf{EDF}(\mathbf{q}, \mathbf{m}(d)) = \left\{ \begin{array}{ll} (d < d_2 - c_2 & \implies [m_1(d_1, c_1), \underline{m(d, c)}, m_2(d_2, c_2), \dots, m_k(d_k, c_k)]), \\ (d_2 - c_2 \leq d < d_3 - c_3 & \implies [m_1(d_1, c_1), m_2(d_2, c_2), \underline{m(d, c)}, \dots, m_k(d_k, c_k)]), \\ \dots & \\ (d_{k-1} - c_{k-1} \leq d < d_k - c_k & \implies [m_1(d_1, c_1), m_2(d_2, c_2), \dots, \underline{m(d, c)}, m_k(d_k, c_k)]), \\ (d_k - c_k \leq d & \implies [m_1(d_1, c_1), m_2(d_2, c_2), \dots, m_k(d_k, c_k), \underline{m(d, c)}]) \end{array} \right\}$$

where  $c$  is a fresh clock not assigned to any message in  $q$ .

Fig. 2. Scheduling functions acting on a queue  $q = [m_1(d_1, c_1), \dots, m_k(d_k, c_k)]$

pendix. Examples of such functions are given in Fig. 2. For simplicity we use the notation  $g \Rightarrow q$  to mean that once guard  $g$  is satisfied the scheduler should produce the queue  $q$ . In our implementation, we will use a timed automaton to act as both the queue and the scheduler function (cf. Section 3.2). As discussed in [Jaghouri et al. 2009], we only consider non-preemptive schedulers because preemption leads to undecidability.

The system is closed if all output actions of all actors are included in the union of their input actions. The behavior of a closed system is then defined by a timed automaton, called the *system automaton*, in which every location contains the local states  $(l_i, q_i)$  of all actors. Without loss of generality, we assume that the sets of method names for different actors are disjoint. The system may take a transition when one actor can make a move, i.e., actors are interleaved. Actors can progress in the following ways:

- An actor is in local state  $(l_i, q_i)$ . Its running method can move from  $l_i$  to  $l'_i$  with a transition that does not involve sending messages. Then the system automaton can take an invisible transition, i.e., with a  $\tau$  action, by changing the local actor state to  $(l'_i, q_i)$ .
- The transition taken by an actor involves sending a message  $m$  with deadline  $d$ . The scheduling function of the receiving actor is then used to determine all possibilities for inserting this message into its queue (cf. Fig. 2). For each possibility, the system automaton includes a transition labeled  $m(d)$  and with the respective guard. Note that the actual contents of the queue will satisfy exactly one of these guards.
- The currently running method is in a final location. An invisible transition, labeled  $\tau$ , removes the method at the queue head and starts the next method.

The detailed definition of the system behavior is given in the appendix. In the following, we explain how to use UPPAAL to implement actors and behavioral interfaces using an example.

### 3.2. Modeling in UPPAAL: A Peer-To-Peer Case Study

Peer-to-peer systems are a commonly used way of sharing data as well as chat-based communication. Contrasted to a client-server architecture, these systems are called peer-to-peer because all nodes can act both as a server and a client; simply said, they are all peers. We model and analyze a hybrid peer-to-peer architecture (like in Skype or BitTorrent), where a central server (called the *broker* or *tracker*) keeps track of all active nodes in the system<sup>1</sup>.

<sup>1</sup>These UPPAAL models are available at <http://www.cwi.nl/~jaghouri/P2P/>.

To start communication, a node acts as a client and asks the broker to connect it to another node. In case of Skype, the client provides the Skype ID to the broker. In a file-sharing system, a keyword is provided in order to search for some data, for example, the name of a song. The broker connects the client to a proper server, e.g., with the given Skype ID, or having the song with the given name. The two nodes then communicate directly by sending requests and replies.

Each node, upon creation, registers its ID/data with the broker. In the case of a file-sharing system, the nodes may obtain new data after every round of communication with other nodes. In this case, they need to update their information registered at the broker.

In UPPAAL, we model for each actor three parts: the behavioral interface, the methods and the scheduler (which in turn include a queue). These automata are parameterized on the identity of the actor itself (written as `self`), and the identifiers of the actors communicating with it (called its known actors). In this case, the known actor of a peer is a `broker`, and a broker has some `Peers` as its known actors. To have more than one instance of an actor, we instantiate the scheduler and method automata and provide different identity values (i.e., `self`) to different actor instances.

*Communication.* We use the channels `invoke` and `delegate` for sending messages. The channel `invoke` has three dimensions (parameters), the message name, the sender and the receiver, e.g., `invoke[connect][Peer][self]!`. This way, actors instantiated from the same automata will have disjoint method names by assigning different identities to their `self` parameter. By setting both sender and receiver as `self` (in method automata), one can invoke a self call (when a deadline is to be given, as explained next). The `delegate` channel is used for delegation. The self call made using the `delegate` channel inherits the deadline of the currently running method (it is taken care of by the scheduler automaton). Since a delegation is used only for self calls, no sender is specified (it has only two parameters).

*Deadlines and Parameters.* We take advantage of the fact that when two edges synchronize, UPPAAL performs the updates on the emitter before the receiver. Hence we can use global variables for passing information. In this case study, we use variables `deadline` and `srv` to pass deadlines and the parameter to `SReq` message (cf. Section 3.2.2), respectively. The emitter sets the desired value into the corresponding variable which is read by the receiver. The receiver, however, cannot use this value in its guard, as guards are evaluated before updates. We define these variables as *meta*, i.e., they are not kept in the state, which implies that their values must be stored properly by the receiver.

*3.2.1. Behavioral interfaces.* The first thing to model for an actor is its behavioral interface. Following the explanation in previous subsection, we model a behavioral interface to represent the environment to the actor. To enable synchronization between outputs of method automata and output actions of the behavioral interface (and similarly between inputs of the scheduler and inputs of the behavioral interface), we use the `!` sign for inputs and `?` for outputs of the behavioral interface.

Both the broker and peers have relatively independent behavior on their server side and client side. Therefore, we model these two sides using independent automata in Fig. 3, which need to be interleaved in order to produce the complete behavior of the actor. Since the messages sent by different peers to the broker are also independent, the client and server side automata of the broker are defined per peer. Fig. 3 shows the messages that a broker object may use to communicate with one peer.

On the server side, the broker specifies only that a peer must register its local information once. However, on the client side, the broker expects to receive requests (`CReq` messages) from a peer repeatedly. For each request, the broker connects it to a server, i.e., the ID of the server is sent as a parameter of the `connect` message to the client; outgoing parameters are not captured in the behavioral interface. We assume that a request by a client is always

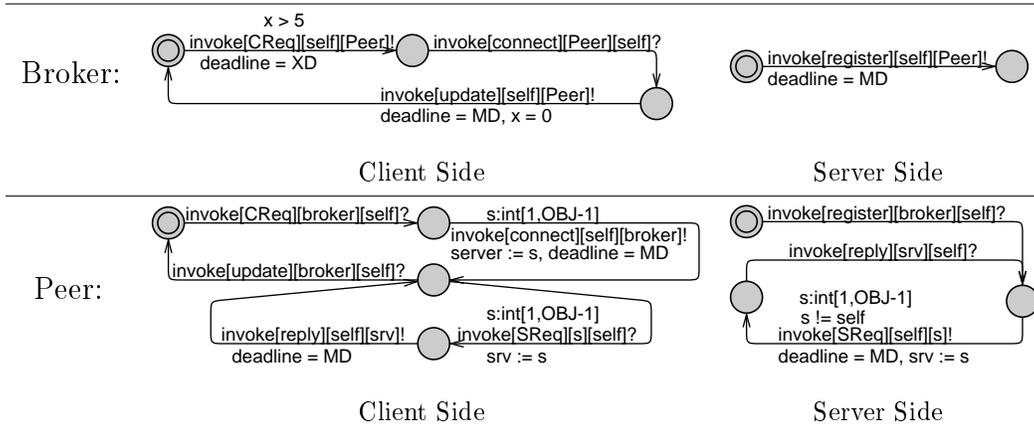


Fig. 3. Behavioral interfaces for Broker and Peer as interleaving of client-side and server-side automata

successful, i.e., every data item searched for is available. The connections between peers is transparent to the broker. Assuming that the client peer has obtained new data after this connection, it should update its registry at the broker (because it can now provide more data on its server side).

Similarly, the server side behavioral interface of a peer starts by registering its data with the broker to initialize its operation. Then it can receive requests (`SReq` messages) and send replies to other peers. We opt for a simple scenario, i.e., each server or client handles only one request at a time. The peer may accept an `SReq` message from any peer ( $s : \text{int}[1, \text{OBJ}-1]$ ) excluding itself ( $s \neq \text{self}$ ). It may only send a `reply` message to the same peer; this is ensured by means of the `srv` variable.

The behavioral interface of the peer is similar to broker on the client side, too, except that it additionally models the communication with a server after a connection has been established. The client can send any number of requests per connection, although only one at a time. Furthermore, the incoming parameter of the `connect` message is also captured with a select expression  $s : \text{int}[1, \text{OBJ}-1]$  in UPPAAL, which means that it may receive the ID of any peer. The global variable `server` is used for communicating this parameter (like the `deadline` variable).

**3.2.2. Broker and Peer Actors.** Fig. 4 shows the method automata of the broker and peer. In this implementation, each method is modeled as a separate automaton. A method may start its behavior when it receives a signal on the `start` channel from the scheduler. After accomplishing its tasks, it sends a signal on the `finish` channel to the scheduler, who will select the next method for execution (see next subsection).

The `initial`, `register` and `update` methods take one time unit to execute. These methods do not perform any computation as we abstract from the data. The `CRReq` method nondeterministically selects a server and sends a `connect` message back to the sender. The variable `sender` is set by the scheduler to refer to the sender of a message. The ID of the selected server is sent using the `server` variable.

In addition to `initial`, a peer implements the `connect` and `reply` methods as a client, and the `SReq` method as a server. Furthermore, the method `userReq` simulates a user who initiates a search request by sending a `CRReq` message to the broker. The `userReq` message is sent first by the `initial` method and then by the `reply` method in order to create a loop. Notice that this implementation of a peer sends exactly one request per connection, while the behavioral interface allows for any number of requests.

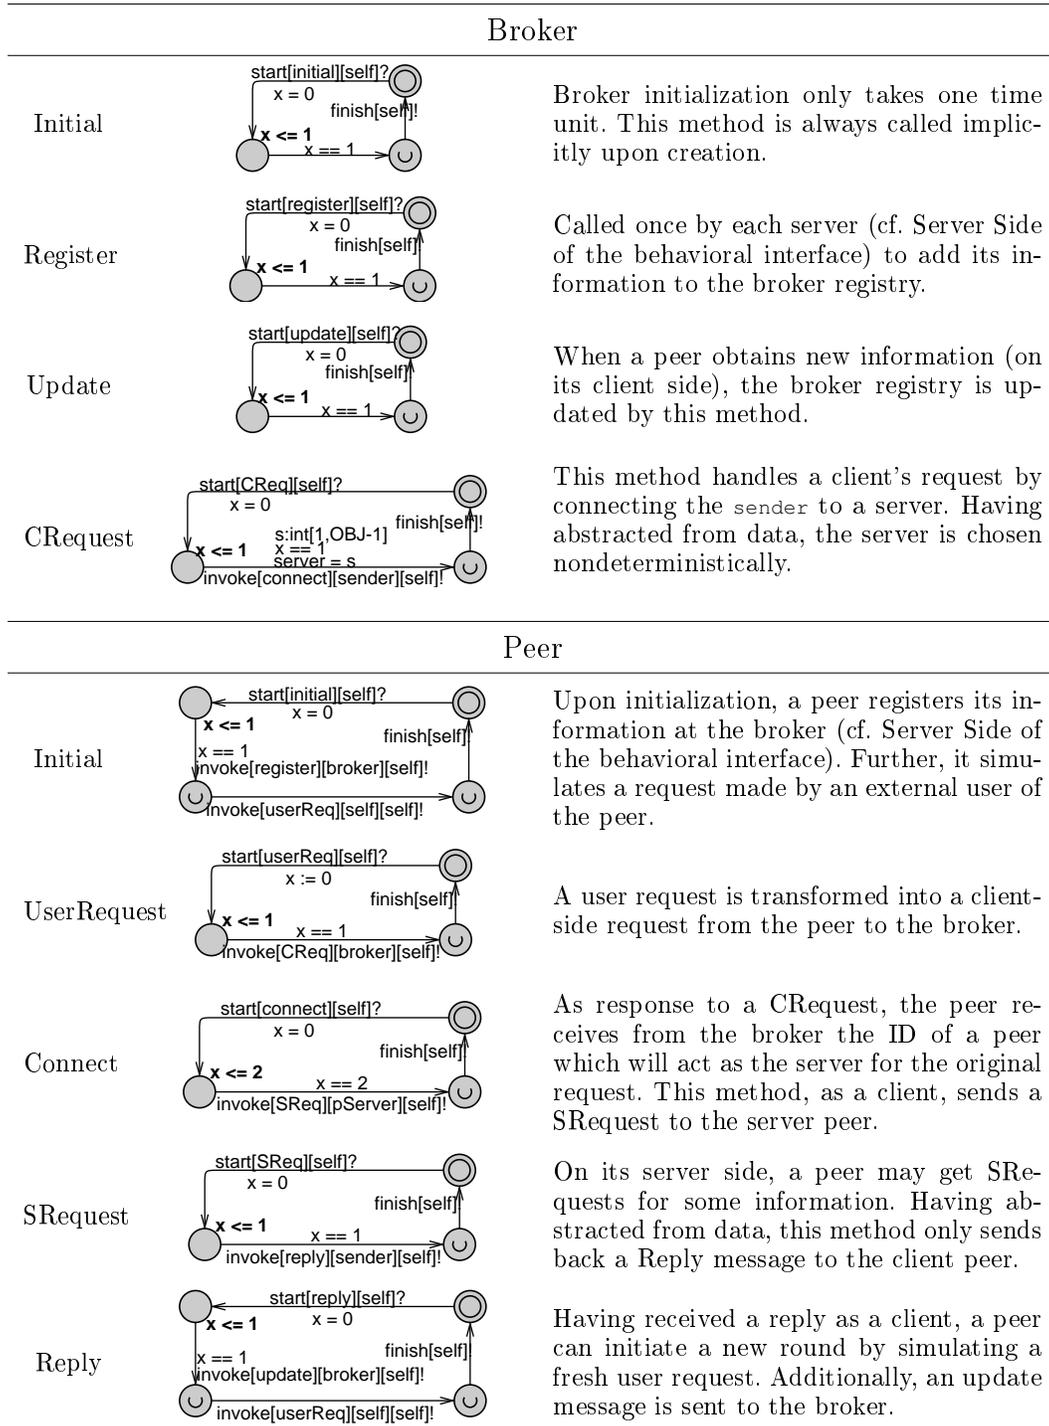


Fig. 4. Method automata defining the Broker and Peer actors

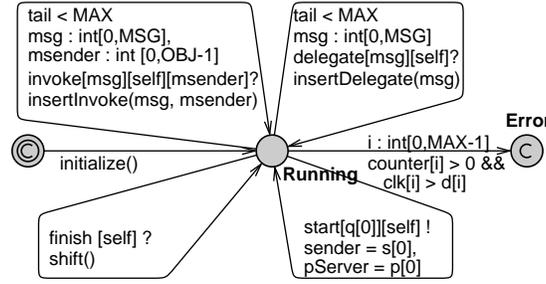


Fig. 5. A first-come first-served scheduler for the Peer actor.

**3.2.3. Modeling the Scheduler.** The scheduler for each actor, containing also its message queue, is modeled separately as a timed automaton (see Fig. 5). The *queue* is modeled using arrays in UPPAAL and thus it can be modeled compactly, i.e., without different locations for different queue states. The scheduler automaton begins with putting an `initial` message in the queue via the `initialize` function. In this section, we briefly describe the ideas on how to model this automaton whereas a detailed explanation including the function definitions is given in the appendix.

The scheduler is input-enabled, i.e., it can receive any message from any sender on the `invoke` channel. The queue stores along with each message its sender and deadline. A free clock is assigned to each message and reset to zero upon insertion in the queue. These are in the `insertInvoke` function. By reusing this clock, a new message may inherit the remaining deadline of another message; this is captured in the `insertDelegate` function. If a clock `clk[i]` is assigned to a message (`counter[i] > 0`) and passes its deadline (`clk[i] > d[i]`), the scheduler moves to an `Error` location. To check schedulability, `Error` must not be reachable and no queue overflow may occur, i.e.,  $\text{tail} \leq \text{MAX}$  must always hold.

In this implementation, new messages are added to the back of the queue, and finished methods are removed from the queue by shifting. When a method finishes, the *scheduling policy* determines the next method for execution. This can be defined as a guard on the transition with the `finish` channel. To implement a “First Come First Served” (FCFS) scheduler, the scheduler in Fig. 5 always starts the message at the head of the queue (`q[0]`); the sender of this message is at `s[0]`. Recall that the `connect` method can take a parameter `pServer`. This is also taken from the parameters array, i.e., `p[0]`. More complex policies could depend on extra information like priorities or the remaining deadline of each message (e.g., in Earliest Deadline First) as described in the appendix.

#### 4. MODULAR SCHEDULABILITY ANALYSIS

An actor is *schedulable* if and only if it finishes all of its tasks within their deadlines. In principle, actors have infinite queues, but we have shown in [Jaghoori et al. 2009] that schedulable actors do not put more than  $\lceil d_{max}/b_{min} \rceil$  messages in their queues, where  $d_{max}$  is the longest deadline for the messages and  $b_{min}$  is the shortest termination time of its method automata. One can calculate the best case runtime for timed automata as shown by Courcoubetis and Yannakakis [Courcoubetis and Yannakakis 1992]. This is important because finite queues make it possible to use model checking techniques for schedulability analysis. Formally, this leads us to the following corollary.

**COROLLARY 4.1 (SYSTEM SCHEDULABILITY).** *A closed system of actors (with unbounded queues) is schedulable if and only if none of the actors exceeds the queue limit of  $\lceil d_{max}/b_{min} \rceil$  and no message in a queue misses its deadline.*

Table I. Schedulability analysis of Peer and Broker individually (left) and in a system (right). The number of known actors of a broker changes its behavioral interface and hence its analysis time.

Single Actor	# Known Actors	Time
Peer	1 Broker	0:07
Broker	1 Peer	0:00
	2 Peers	0:02
	3 Peers	1:34:38
	4 Peers	$\infty$

Composed System	Time
1 Peer & 1 Broker	0:00
2 Peers & 1 Broker	4:02
3 Peers & 1 Broker	$\infty$

We can use the value  $\lceil d_{max}/b_{min} \rceil$  for bounding the queues of schedulers as explained in the previous section. These automata have a special location `ERROR` such that a missed deadline results in an error. The system is then schedulable if the `ERROR` location is not reachable and no queue overflow occurs. In other words, schedulability analysis is reduced to reachability analysis in a tool like UPPAAL and thus it is decidable. However, the intrinsic asynchrony of actors and their message buffers will lead to state space explosion for larger systems. This can be avoided by modular analysis of the actors. To this end, we use the behavioral interface of every actor as a *contract* between the actor and its environment. Below, we describe how to check whether, firstly the actor itself, and secondly the environment in which it is used, respect this contract. We refer to the latter as compatibility check.

#### 4.1. Individual Actor Analysis

Analyzing actors in isolation is hindered by the fact that the methods of an actor can in theory be called in infinitely many ways. However, taking the behavioral interface as the contract to which the actor should adhere, it is reasonable to restrict only to the incoming method calls specified in its behavioral interface. In other words, we use the behavioral interface as a driver where the input actions correspond to the incoming messages. Incoming messages are buffered in the actor; this can be interpreted as creating a new task for handling that message. The behavioral interface doesn't capture the internal tasks triggered by self calls. Therefore, one needs to consider both the internal tasks and the tasks triggered by the behavioral interface, which abstractly models the acceptable environments.

The semantics of an isolated actor can be defined as an *isolated actor automaton*. The states of this automaton are written as  $(l_i, q_i, b_i)$  where  $l_i$  shows the current location of the currently running method,  $q_i$  reflects the contents of the actor queue, and  $b_i$  is the current location of the behavioral interface; a full characterization is given in [Jaghoori et al. 2009]. Similar to the case of a complete system, schedulable isolated actors do not put more than  $\lceil d_{max}/b_{min} \rceil$  messages in their queues.

**COROLLARY 4.2 (ACTOR SCHEDULABILITY).** *An actor is schedulable with respect to its behavioral interface if and only if it does not exceed the queue limit of  $\lceil d_{max}/b_{min} \rceil$  and no message in the queue misses its deadline.*

*Peers and Broker.* To analyze the broker, we need to know from how many peers it may receive requests. The automata representing the behavioral interfaces of the broker (cf. Fig. 3) need to be replicated for every peer, and these instances should be interleaved. For more efficient analysis, queue sizes smaller than  $\lceil d_{max}/b_{min} \rceil$  can be tried first; and, to handle three peers it turns out that the broker can manage with a queue of size 7. For bigger systems, the model checking becomes very time consuming and intractable. To improve efficiency one can follow the guidelines in Credo methodology [Grabe et al. 2009]. The analysis of each peer can be performed with one instance of its client side and server side behavioral interfaces.

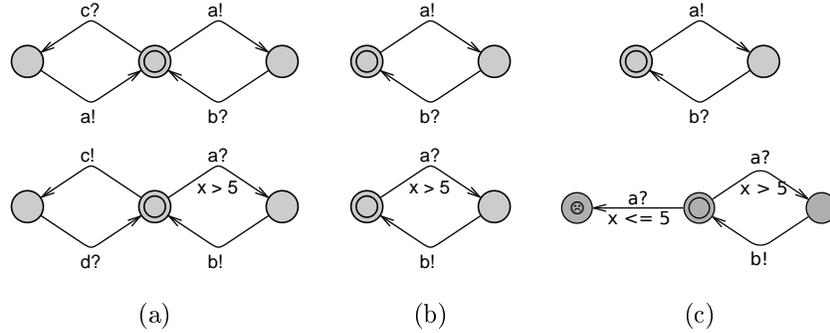


Fig. 6. Compatibility cannot be checked at the level of behavioral interfaces.

Table I summarized schedulability analysis times of different configurations. The table on the left shows the analysis of individual actors. The high level of asynchrony introduced by increasing the number of known actors for the broker results in highly exponential increase in the analysis time. For comparison, the table on the right shows how long it would take to analyze a complete system. We can see that analysis of a complete system takes a much longer time and becomes intractable faster than individual actors. The combination of model checking and testing for modular schedulability analysis proposed in this paper is an effective method to overcome this problem.

#### 4.2. Compatibility check

Once an actor is proved to be schedulable with respect to its behavioral interface, it can be used as an off-the-shelf component. A system composed of individually schedulable actors is itself schedulable if the actual use of the actors in this system is compatible with their behavioral interfaces. For each actor, the behavioral interface abstractly models its observable behavior in terms of the messages it may receive and the messages it sends.

Ideally, it would be enough to check compatibility only considering the behavioral interfaces, for example, by checking deadlock freedom in the composition of the behavioral interfaces. Unfortunately, this is not possible in a real-time system. We explain this using three sample pairs of behavioral interfaces shown in Fig. 6. On the one hand, behavioral interfaces cannot be used to disprove compatibility. Consider the two automata in Fig. 6.(a). If both automata take their left transitions, i.e., communicate by message  $c$ , there will be a deadlock because of the mismatching provided and required messages. However since these behavioral interfaces are abstractions of object behaviors, such a mismatch could be due to a spurious behavior not possible in the real system model. In other word, the implementations of these behavioral interfaces could happily communicate only  $a$  and  $b$  messages.

On the other hand, behavioral interfaces cannot be used to prove compatibility, either. For example, the automata in Fig. 6.(b) can be composed with no problem, e.g., no deadlock occurs. Such a compatibility, as for example defined for timed interfaces [de Alfaro et al. 2002], means that compatible implementations exists; but this does not guarantee compatibility of every possible implementation. An actor implementing the top interface may send  $a$  outside the time constraint required by the bottom interface. In general, a behavioral interface does not reflect the precise timing of the send action by the real system model. In the work of [David et al. 2010], this problem is avoided by requiring the specifications to be input-enabled, like Fig. 6.(c) where unacceptable inputs leads to an error location. This is, however, too restrictive because for example, it makes the example in Fig. 6.(c) incompatible whereas we know already that compatible implementations exist.

The only solution to this problem is to take both the system composition and the behavioral interfaces into account. Intuitively, a running system of actors is compatible with their

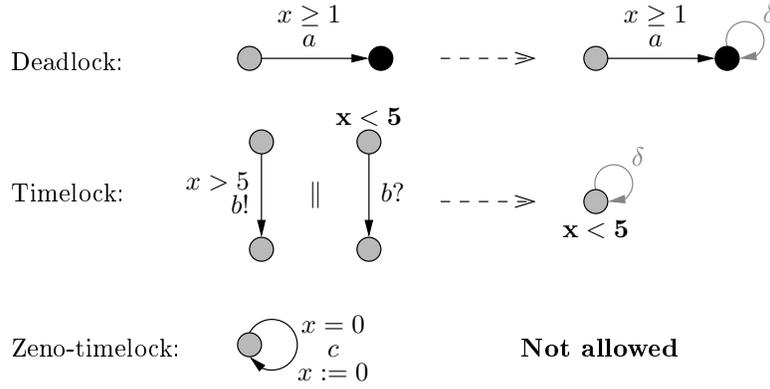


Fig. 7. Construction of the suspension automaton

behavioral interfaces if its observable behavior is captured by the composition of the behavioral interfaces of the participating actors. Checking compatibility is prone to state-space explosion due to the size of the system; we avoid this by means of the testing technique described in Section 5. We first formally define compatibility in terms of a refinement relation.

*Quiescence in refinement.* In the context of timed automata, an observable behavior is either an observable action (any action except  $\tau$ ), the passage of time (a delay) or blocking (also called quiescence, i.e., the absence of observable actions). Actions and delays are already taken into account in timed traces, which describe the possible distribution of observable actions in time. To be able to take blocking into account, we need to represent them explicitly in the specification. There are three different scenarios for blocking a real-time system:

*Deadlock.* No action is possible but time can go on.

*Timelock.* Time is stopped, no action and no delay is possible (in the left hand side of Fig. 7 the synchronization cannot happen due to the mismatching invariant and guard).

*Zeno-timelock.* Infinitely many actions can occur in finite time.

Deadlocks and timelocks can occur as a result of composing the behavioral interfaces, and are therefore allowed in our refinement definition; the specifications, however, should not include zeno-timelocks. For a correct refinement, the system may deadlock (resp. timelock) only if the composition of behavioral interfaces deadlocks (resp. timelocks). Locations with a deadlock or timelock are called *quiescent*. To explicitly specify quiescence in the specification, we add a loop on each blocking location labeled by a new action  $\delta$ , considered to be an observable action [Tretmans 1996] (see the right hand side of Fig. 7). The automaton obtained by adding  $\delta$  actions is called a *suspension automaton*.

*Definition 4.3 (Suspension automaton).* Let  $A = (L, l_0, E, I)$  be a timed automaton over clocks  $C$  and actions  $Act$ . The *suspension automaton* of  $A$  is the timed automaton  $\Delta(A) = (L, l_0, E_\Delta, I)$  over  $C$  and  $Act \cup \{\delta\}$  where  $E_\Delta$  is defined as follows:

$$E_\Delta = E \cup \{l \xrightarrow{\delta} l \mid l \text{ is quiescent}\}$$

The traces of the suspension automaton, called suspension traces, then represent all the observable behaviors of this automaton. The set of all suspension traces of a timed automaton  $A$  is the set  $Traces(\Delta(A))$  denoted by  $STraces(A)$ .

In our models,  $S$  represents the system composition, and  $B$  is the synchronous product of the behavioral interfaces. Thus  $B$  and  $S$  represent two timed automata on the same set of actions, while  $S$  also contains  $\tau$  actions. In the definition of refinement below, we consider observable suspension traces of  $S$  and all suspension traces of  $B$ .

Further we need to consider deadlines in refinement. Recall that in the behavioral interfaces, only *input* actions are assigned deadline, whereas in the system of actors, the *output* actions in method automata have deadlines. Input actions in  $B$  provide the guaranteed deadlines (checked during individual actor analysis) whereas output actions in methods specify the required deadlines. In the definition of refinement below, we define that a deadline required by an action in  $S$  may not be smaller than the deadline guaranteed by a matching action in  $B$ .

*Definition 4.4 (Refinement).*  $S$  is a *refinement* of  $B$ , denoted by  $S \sqsubseteq B$ , if and only if for every observable trace  $\sigma = \dots, t_i, m_i(d_i), \dots$  in  $S\text{Traces}_{obs}(S)$ , there exists a matching trace  $\sigma' = \dots, t_i, m_i(d'_i), \dots$  in  $S\text{Traces}(B)$ , such that  $d_i \geq d'_i$  for every  $i$ .

*Definition 4.5 (Compatibility).* The system  $S$  is compatible with the behavioral interfaces if and only if  $S \sqsubseteq B$  where  $B$  is the synchronous product of the behavioral interfaces.

The actors used in a system are proved individually schedulable with respect to their behavioral interfaces (as explained in Section 4.1). The following theorem proves that compatibility implies schedulability of the whole system provided that individual behavioral object models are schedulable. Intuitively, this means that every message in the system will be finished within the designated deadline.

**THEOREM 4.6 (SYSTEM SCHEDULABILITY).** *A system composed of a set of actors  $O_1, \dots, O_n$  is schedulable, if every actor  $O_i$  is individually schedulable and the system is compatible with the behavioral interfaces of the actors.*

**PROOF.** To prove this we can assume that the system is compatible but not schedulable. This means that there is a trace  $\delta = t_1, a_1, t_2, \dots, a_k, t_{k+1}$  of the system automaton (cf. Appendix A) in which one of the actors, say  $O_j$ , drives the system to the *Error* state, i.e., either the queue of  $O_j$  is overflowed or a task in its queue misses its deadline. We show that this requires the existence of a trace in  $B_j$  that drives  $O_j$  to the *Error* state, which contradicts the schedulability assumption.

Due to compatibility,  $\delta_{obs}$  exists in the product of behavioral interfaces. This trace can be projected onto the behavioral interface of  $O_j$  alone by removing the delay-actions  $t_i, a_i$  for every  $a_i$  that is not in the action set of the behavioral interface of  $O_j$ . Call the resulting trace  $\delta_{obs}|_j$ . We can compute a set of traces in ‘isolated actor automaton’ of  $O_j$  (cf. Section 3) as  $T = \{\varphi_i : \varphi_{i_{obs}} = \delta_{obs}|_j\}$ . Since the actor individually is schedulable, these traces do not lead to *Error* state, i.e., given this sequence of inputs and outputs, none of the tasks in the queue of  $O_j$  misses its deadline, nor a queue overflow occurs.

On the other hand, the trace  $\delta$  corresponds to a run of the system:

$$(\{l_0^1, \dots, l_0^n\}, \mathbf{0}) \xrightarrow{a_1} \dots \xrightarrow{a_{k-1}} (\{l_{k-1}^1, \dots, l_{k-1}^n\}, u_{k-1}) \xrightarrow{a_k} (Error, u_k)$$

where  $l_i^j$  is the location of actor  $j$  after  $i$  steps. Formally,  $l_i^j = (s, Q)$  where  $s$  is the location of its currently running task and  $Q$  is the current task queue. For brevity the delay transitions are not shown. We project the trace  $\delta = t_1, a_1, t_2, \dots, a_k, t_{k+1}$  onto the actions of  $O_j$ , by removing the delay-actions  $t_i, a_i$  such that  $l_{i-1}^j = l_i^j$ . We represent the resulting trace as  $\delta|_j = u_1, b_1, u_2, \dots, b_h, u_{h+1}$ .

By considering the definition of system automaton, we can show that  $\delta|_j \in T$ . This requires that  $\delta|_j$  drives  $O_j$  to the *Error* state, which is in contradiction with the schedulability assumption.  $\square$

Since  $B$  is deterministic, checking trace inclusion becomes decidable [Alur and Dill 1994; Simons and Stoelinga 2001], but due to the size of  $S$ , it may be susceptible to state-space explosion. To avoid this, we propose a method for *testing* trace inclusion in the next section. In particular, we want to be able to exhibit a counter-example if some incompatibility is found.

## 5. COUNTER-EXAMPLE ORIENTED TESTING: COMPATIBILITY CHECK

We will show in this section how compatibility, defined as a refinement relation based on trace inclusion, can be tested. Trace inclusion is a usual notion of correctness (or conformance) between a system and its abstract specification  $B$  in formal testing frameworks [Hessel et al. 2008]. A naive approach involves taking a trace from the system model and check if it is a valid trace in the abstract specification. This is not practical because the system model is very big and is not a suitable source of generating test cases. Our idea is to generate test cases based on traces from  $B$  and use it to restrict the system behavior. As long as the system can follow this trace, the test case looks for possible violations of the refinement relation.

Testing compatibility gives rise to some issues which are not common in standard frameworks. First, the system under test is a model and not a real implementation. We do not take advantage of our knowledge of the model, so it can be seen as black-box testing, but the execution of test cases will be simplified as we can apply tools that can systematically explore a model. Another difference from usual frameworks is that the system involves internal actions that are not specified in the behavioral interfaces. The consequence is that the test, built from the abstract specification only, will not be able to fully control the system under test during its execution. This leads to a lot of non-determinism but this can be solved by using a model checker to execute a test case. Last but not least, our main goal is to find a counter-example in the case of wrong refinement. Then test cases must be as “rigid” as possible to take any incorrect behavior into account. We will formally define rigidity in this section as a characteristic of counter-example oriented testing. First, we formally define a test case.

We are given two timed automata: one is the system of actors that we call the model under test  $MUT$ , and the other is the product of the behavioral interfaces of these actors. To test compatibility, we take the suspension automaton of the latter as the abstract specification. We denote this by  $B$ , which is then a deterministic timed automaton over the action set  $Act_B$ .  $MUT$  is a timed automaton over the set of actions  $Act_{MUT} = Act_B \cup \{\tau\}$ . A test case is a deterministic timed automaton without loops whose leaves are labeled with verdicts.

*Definition 5.1 (Test case).* Let  $B$  be a timed automaton over  $Act_B$ . A *test case* for  $B$  is a deterministic acyclic timed automaton  $TC = (L, l_0, E, I)$  over  $Act_B \cup \{\tau\}$ , in which all leaf locations (i.e., those with no outgoing transitions) are labeled with a verdict **Pass** or **Fail**. We refer to a set of test cases as a test set.

A verdict labeling a location allows us to evaluate an execution of the test case terminating on this location. The **Pass** verdict is reachable via only one path, which covers exactly the intended behavior we are testing for. This means that the system fulfilled the test case requirements. To find a counter-example to refinement, we need to search for locations marked **Fail**. These are the locations that are reachable with forbidden behaviors of the system (a non-specified action or an action happening outside its time constraints in the specification of  $B$ , for example). If the system deviates from the behavior aimed at by the test case without violating refinement, the test may terminate prematurely resulting in an inconclusive verdict. This is similar to the ‘timed failures’ used in [Reed and Roscoe 1999] as a semantic model for timed CSP.

Recall that proving compatibility implies schedulability of the system. However, violating refinement and thus compatibility does not per se imply the violation of schedulability.

*intuition: untimed →  
timed; why  
determinism*

Nevertheless, considering the assume-guarantee approach, it does violate the assumptions on schedulability of individual actors specified in the behavioral interfaces. Therefore, by means of testing one can find and remove counter-examples to compatibility and, as a result, attain more confidence in schedulability of the system.

The fact that a test case is deterministic means that from any state  $l$ , for any action  $a \in Act_B$ , all transitions from  $l$  labeled by  $a$ , as well as the transition labeled by  $\tau$  if any, have disjoint guards: for all  $a \in Act_B$ , for all transitions  $l \xrightarrow{g_i, a, r_i} l'_i$ ,  $1 \leq i \leq k-1$  and  $l \xrightarrow{g_k, \tau, r_k} l'_k$ , the formula  $\bigwedge_{1 \leq i \leq k} g_i$  is unsatisfiable.

### 5.1. Generating a test case

The idea is to take a timed trace from the abstract specification  $B$  and turn it into a test case. Since the exact timing of actions in a timed trace make the test case too restrictive, we take instead a linear timed automaton  $T$ . This consists of a sequence of transitions from  $B$  representing a set of timed traces, but corresponding to exactly one untimed trace. As shown in Fig. 8,  $T$  contains a sequence of transitions written as  $l_{i-1} \xrightarrow{g_i, m_i, r_i} l_i$ ,  $1 \leq i \leq n$ . Such a sequence abstractly represents a desired system behavior (the test purpose).

The sequence of transitions of  $T$  corresponds to the behavior we want to test so the last location must be labeled **Pass**. All other locations are completed as follows, such that any forbidden behavior makes the test fail. If a location has an invariant  $h_i$  in  $B$ , violating this invariant must make the test fail; thus, a transition labeled with  $\tau$  and with guard  $\neg h_i$  leading to **Fail** is added. Furthermore, no other transition may be taken if the invariant is violated; this is ensured by conjunction of guards of all other transitions with  $h_i$ . Additionally, every behavior which is not allowed in  $B$  is forbidden, so for every action, a transition labeled by this action and whose guard is the complement of all the existing guards for this action leads to a **Fail** location; this guard is computed in  $g_j$ . Any trace leading to **Fail** is an example of behavior not allowed in the abstract behavior specification.

*Example 5.2.* In Fig. 9,  $B$  shows the suspension automaton of an abstract specification,  $T$  is a selected sequence of transitions from  $B$  and finally  $TC$  is the test case generated from  $T$  by the algorithm. The location invariant  $\mathbf{x} \leq \mathbf{10}$  is kept in bold face in the test case only to show its effect on guards. The transition labeled by  $Act$  stands for three transitions labeled by  $a$ ,  $b$  and  $c$  with guard *true*.

The theorem below states that the timed automaton we obtain by this construction is a test case in the sense of Definition 5.1. The proof is straightforward by following the steps in the algorithm.

**THEOREM 5.3.** *Let  $B$  be a timed automaton over  $Act_B$ . For any linear timed automaton from  $B$ , the automaton generated using the algorithm in Fig. 8 is a test case.*

*Remark.* The starting point for test case generation is a sequence of transitions in the specification. Given a desired reachability property  $\varphi$ , we can generate such a sequence of transitions automatically. We start by model-checking  $\varphi$  on the suspension automaton of the specification  $B$ . The diagnostic trace produced by the model-checking tool gives the sequence of moves that have to be made by this automaton and the required clock constraints needed to reach the targeted location. This method is in parts similar to [Hessel et al. 2008]. Instead of checking for a reachability property, one can also use the simulation feature of a model-checker to generate specific hand-made traces. Another interesting property is a deadlock or timelock in the abstract specification  $B$ . Although a correct refinement is theoretically allowed to deadlock or timelock in such cases, too, such situations are in practice undesired. Therefore, such traces could also contribute to good test cases for checking system correctness. Note that a timelock in our models does not violate schedulability because when time stops no deadline is missed, but such a scenario is in fact an unrealistic situation.

**Inputs**  $B = (L_B, l_{0_B}, E_B, I_B)$ : A timed automaton specifying the abstract behavior  
 $T = (L_T, l_0, E_T, I_T)$  such that  $L_T \subseteq L_B$  in linear form as:



**Output**  $TC = (L_T \cup \{f\}, l_0, E_C, TR)$ : A timed automaton representing the test case where  $TR$  is a function that always returns *true* as location invariant

**Constructing the transitions:**

```

 $E_C := \{ \}$ 
for each  $i \in [0 .. n - 1]$  do
   $h_i := I_B(l_i)$  // the invariant of  $l_i$  in  $B$ 
   $E_C := E_C \cup \{l_i \xrightarrow{\neg h_i, \tau, \emptyset} f\}$ 
   $E_C := E_C \cup \{l_i \xrightarrow{g_i \wedge h_i \wedge (d \geq d_i), m_i(d), r_i} l_{i+1}\}$ 
  for each action  $m \in Act_B$  do
     $g_f := false$ 
    for each transition  $l_i \xrightarrow{g, m(d'), r} l' \in E_B$  do  $g_f := g_f \vee (g \wedge d \geq d')$  endfor
     $E_C := E_C \cup \{l_i \xrightarrow{h_i \wedge \neg g_f, m(d), \emptyset} f\}$ 
  endfor
endfor

```

**Verdicts:**

Label  $l_n$  with the verdict **Pass**  
 Label  $f$  with the verdict **Fail**

Fig. 8. Test case generation algorithm

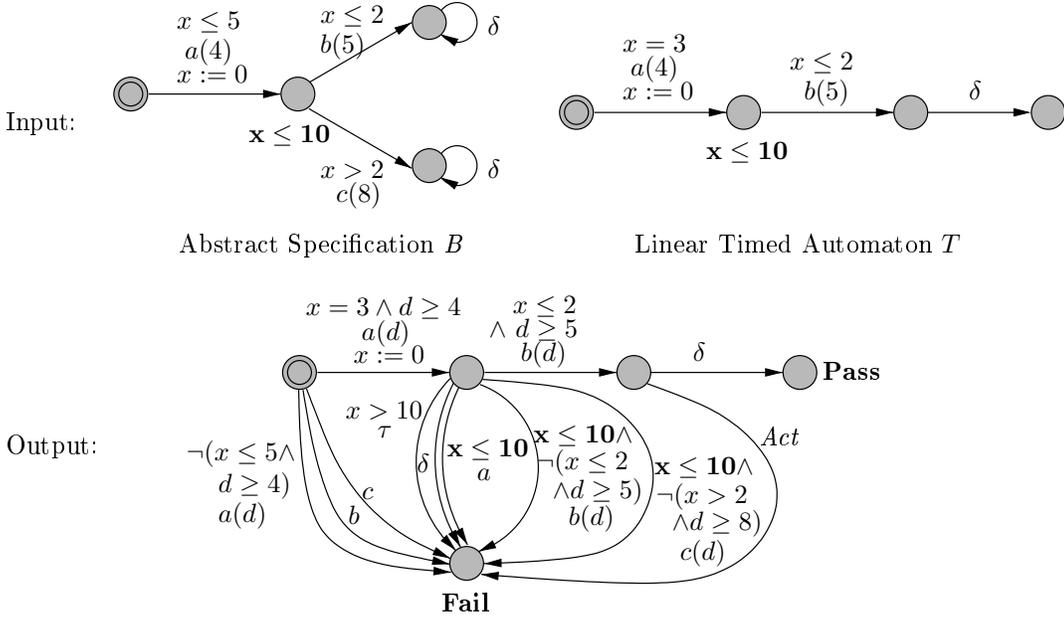


Fig. 9. In this example, the input  $B$  is a suspension automaton. Actions with no deadlines in the output test case imply unrestricted deadline values. Label *Act* means all actions are possible.



5.2.2. *Exhaustiveness.* Soundness is not sufficient to ensure the relevance of test cases. A test case with no **Fail** location is sound but cannot reject any system. We also need to be sure that if the system is a wrong refinement, there exists a test case able to reject it. An exhaustive test set rejects any wrong refinement in the system. In other words, any system that passes the test set is a correct refinement of the specification. A test set is formally defined to be *exhaustive* for the refinement relation  $\sqsubseteq$  if and only if

$$MUT \text{ passes } \mathcal{T} \implies MUT \sqsubseteq B$$

**THEOREM 5.5 (EXHAUSTIVENESS).** *The set of all test cases for  $B$  that can be generated by the algorithm in Fig. 8 is exhaustive for  $\sqsubseteq$ .*

**PROOF.** We must prove that if the system is not a refinement, there exists a test case that makes the system fail. In other words, assuming that there exists an observable suspension timed trace of  $MUT$  not belonging to suspension timed traces of  $B$ , we must show that there exists a test case  $TC$  such that the **Fail** location is reachable in the product  $TC \parallel MUT$ .

Without loss of generality, we consider  $\sigma = t_1 a_1 \dots t_k a_k \in STraces_{obs}(MUT)$ . The corresponding observable run in  $MUT$  is the following:

$$(l'_0, \mathbf{0}) \xrightarrow{d_1} (l'_0, u'_0) \xrightarrow{a_1} (l'_1, u_1) \rightarrow \dots \xrightarrow{d_k} (l'_{k-1}, u'_{k-1}) \xrightarrow{a_k} (l'_k, u_k)$$

If  $\sigma$  is not a trace of  $B$ , then two cases are possible depending on the first behavior diverging from the trace of  $B$ :

- (1) There exists  $i$ ,  $0 \leq i \leq k-1$  such that  $t_1 a_1 \dots t_i a_i \in STraces(B)$  and  $t_1 a_1 \dots t_i a_i t_{i+1} \notin STraces(B)$ , i.e., a delay of  $d_{i+1} = t_{i+1} - t_i$  is not possible in  $B$  at state  $(l_i, u_i)$ . It means that location  $l_i$  has an invariant  $h$  that is violated by  $u_i + d_{i+1}$ . Let  $T$  be a sequence of  $i+1$  transitions from  $B$  such that  $t_1 a_1 \dots t_i a_i \in STraces(T)$ . Let  $TC$  be the test case generated from  $T$  by the algorithm. Since location  $l_i$  has an invariant  $h$ , there is a transition in  $TC$  from  $l_i$  to location  $f$  whose guard is  $\neg h$  and labeled with  $\tau$ . As  $u_i + d_{i+1}$  does not satisfy  $h$ , location  $f$  is reachable in the product  $TC \parallel MUT$  with the trace  $t_1 a_1 \dots t_i a_i t_{i+1} \tau$  corresponding to the run

$$((l_0, l'_0), \mathbf{0}) \xrightarrow{d_1} ((l_0, l'_0), u'_0) \xrightarrow{a_1} \dots \xrightarrow{d_{i+1}} ((l_i, l'_i), u'_i) \xrightarrow{\tau} ((f, l'_{i+1}), u_{i+1})$$

- (2) There exists  $i$ ,  $0 \leq i \leq k-1$  such that  $t_1 a_1 \dots t_i a_i t_{i+1} \in STraces(B)$  and  $t_1 a_1 \dots t_i a_i t_{i+1} a_{i+1} \notin STraces(B)$ , i.e., the action  $a_{i+1}$  is not allowed in  $B$  at time  $t_{i+1}$ . It means that there is no transition in  $B$  from location  $l_i$  labeled with  $a_{i+1}$  whose guard is satisfied by  $u_i + d_{i+1}$ . Let  $T$  be a sequence of  $i+1$  transitions from  $B$  such that  $t_1 a_1 \dots t_i a_i t_{i+1} \in STraces(T)$ . Let  $TC$  be the test case generated from  $T$  by the algorithm. By construction of  $TC$ , there is a transition from location  $l_i$  to location  $f$  labeled with  $a_{i+1}$  whose guard is the complement of all other guards of transitions from  $l_i$  labeled with  $a_{i+1}$ , let us call it  $g$ . Since  $u_i + d_{i+1}$  does not satisfy any of these guards, it satisfies  $g$ . Then location  $f$  is reachable in the product  $TC \parallel MUT$  with the trace  $t_1 a_1 \dots t_i a_i t_{i+1} a_{i+1}$  corresponding to the run

$$((l_0, l'_0), \mathbf{0}) \xrightarrow{d_1} ((l_0, l'_0), u'_0) \xrightarrow{a_1} \dots \xrightarrow{d_{i+1}} ((l_i, l'_i), u'_i) \xrightarrow{a_{i+1}} ((f, l'_{i+1}), u_{i+1})$$

Therefore, the set of all test cases generated by the algorithm is exhaustive for  $\sqsubseteq$ .  $\square$

A sound and exhaustive test set is called *complete*. Completeness is in general impossible to reach, since it usually needs an infinite test set. Thus we know that we cannot in practice find all counter examples to refinement. However, we still want to ensure a certain quality to test cases. For instance, we want to avoid useless sound test cases where all paths lead to **Pass**. Below, we introduce a more practical property for test cases.

**5.2.3. Rigidity.** We are interested in test cases that reject models which behave in a wrong way *along the test case*: the test case should not say **Pass** if it is possible to detect something wrong during the test case execution. We show that any test case generated by our algorithm can detect every wrong behavior occurring along it. We can actually show that we can provide a counter-example for any incorrect refinement occurring along the sequence of transitions the test case is built from. Given a trace  $\sigma \in STraces(B)$ , suppose the action or delay  $e \in Act \cup \{\delta\} \cup \mathbb{R}_+$  is allowed after  $\sigma$  in  $MUT$  but not in  $B$ , i.e.,  $\sigma.e \in STraces_{obs}(MUT) \setminus STraces(B)$ . A test case  $TC$  is rigid for the refinement relation  $\sqsubseteq$  if and only if it rejects any incorrect refinement along the traces of the test case:

$$\sigma \in Traces(TC) \wedge l_0 \xrightarrow{\sigma} l_i, 1 \leq i < n \Rightarrow \sigma.e \in Traces(TC) \wedge l_0 \xrightarrow{\sigma.e} \mathbf{Fail}$$

Intuitively, if  $\sigma$  ends in a non-leaf location in  $TC$ , the test case  $TC$  will observe any one-step divergence after  $\sigma$ . This notion is close to the notions of non-laxness in the untimed setting [Jard et al. 2000] and of strictness in the timed setting [Krichen and Tripakis 2004] but it is stronger. These notions state that if the system behaves in a non-conforming way during the execution of the test case, it must be rejected. Also in our framework, every detected divergence leads to the rejection of the system, but we can add that every divergence is actually detected. This result directly follows from the construction of the test case.

**THEOREM 5.6 (RIGIDNESS).** *Let  $B$  be a deterministic timed automaton and  $T$  be a linear timed automaton built from a sequence of transitions in  $B$ . The test case for  $B$  generated from  $T$  by the algorithm in Fig. 8 is rigid for  $\sqsubseteq$ .*

**PROOF.** We show that for every trace  $\sigma$  of the test case  $TC$  ending in a non-leaf location,  $\sigma.e$  is a trace of  $TC$  leading to **Fail** if  $\sigma.e$  is a trace of  $MUT$  and not of  $B$ .

If  $e \in Act \cup \{\delta\}$ , let  $\sigma = t_1 a_1 \dots t_k \in Traces(TC)$  corresponding to the run

$$(l_0, \mathbf{0}) \xrightarrow{d_1} (l_0, u'_0) \xrightarrow{a_1} (l_1, u_1) \rightarrow \dots \xrightarrow{d_k} (l_{k-1}, u'_{k-1})$$

where  $k-1 \neq n$ . Since  $\sigma.e$  is not a trace of  $B$ , it means that  $e$  is not allowed in  $B$  at location  $l_{k-1}$  after a delay of  $d_k$ . Then, by construction of the test case  $TC$ , there is a transition from  $l_{k-1}$  to the **Fail** location labeled with action  $e$  and whose guard satisfies  $u'_{k-1}$ . Then  $\sigma.e$  is a trace of  $TC$  and  $l_0 \xrightarrow{\sigma.e} \mathbf{Fail}$ .

If  $e \in \mathbb{R}_+$ , let  $\sigma = t_1 a_1 \dots a_k \in Traces(TC)$  corresponding to the run

$$(l_0, \mathbf{0}) \xrightarrow{d_1} (l_0, u'_0) \xrightarrow{a_1} (l_1, u_1) \rightarrow \dots \xrightarrow{a_k} (l_k, u_k)$$

where  $k \neq n$ . Since  $\sigma.e$  is not a trace of  $B$ , it means that a delay of  $e$  is not allowed in  $B$  at location  $l_k$ , due to an invariant  $h$  at this location in  $B$ . Then, by construction of the test case  $TC$ , there is a transition from  $l_k$  to the **Fail** location labeled with  $\tau$  and whose guard  $\neg h$  satisfies  $u_k$ . Then  $\sigma.e$  is a trace of  $TC$  and  $l_0 \xrightarrow{\sigma.e} \mathbf{Fail}$ .  $\square$

### 5.3. Executing test cases in UPPAAL

Recall that we are testing the inclusion of the observable traces of a system  $S$  in the traces of a specification  $B$ . Test cases generated from  $B$  are used to restrict the system behavior and at the same time detect any violations of refinement along the test case. When checking compatibility,  $B$  corresponds to the composition of all behavioral interfaces; if  $S$  is a refinement of  $B$ , the schedulability of  $S$  is guaranteed based on Theorem 4.6.

Our model of a system consists of a set of communicating actors. Each actor, in turn, consists of methods and a scheduler automaton (the scheduler contains a queue). Therefore, a system is a network of timed automata (representing methods and schedulers of all actors) which can be run in UPPAAL. All actions except communication between actors are considered internal.

When submitting a test case, we require that any communication between two actors should synchronize with the test case, as well. Practically, this means that the sender actor (in one of its methods), the receiver actor (in its scheduler) and the test case should synchronize. Since we do not want to change the specification of the model under test, we solve the problem of three-way synchronization by splitting every action in the test case into two steps. At the first step, the sender actor synchronizes with the test case, and *immediately* afterwards, the test case synchronizes with the receiver actor. The urgency between these two steps is modeled by using a ‘committed’ location in the test case between these two steps. For the test case to be able to intercept the messages, we bind all known actors to refer to the test case (see Fig. 10).

A test case is a deterministic automaton and intuitively represents a specific order of exchanging messages between the actors. Although requiring the system to synchronize with the test case resolves part of the non-determinism in the system, the final model is not yet completely deterministic. For example, UPPAAL will still consider the interleaving of the internal actions of different actors. The actor implementations may also contain some internal choices that are not controlled by the test case. In principle, this would mean that a test case needs to be repeated several times and a coverage of different nondeterministic choices requires extra control over the system behavior. To avoid this problem, we take advantage of the fact that we are testing a model and not a real implementation. This simplifies the submission of test cases by means of a tool like UPPAAL as explained in the sequel.

What is important in the execution of a test case is the final verdict. Reaching a **Pass** or **Fail** verdict can in fact be formulated as a reachability property in UPPAAL. Using the model-checker of UPPAAL to execute the test case, one can look for the reachability of the **Pass** or **Fail** location. This way each test case needs to be submitted once. The UPPAAL model checker can provide a diagnostic trace whenever the searched verdict is reachable. A trace leading to the **Fail** verdict shows exactly how and when the system is not compatible.

Using the model checker in this scenario is plausible because the system behavior is controlled by the test case, while model-checking the whole system may not be tractable. We thus avoid state space explosion by restricting verification to the part of system behavior that follows the main line of the test case.

#### 5.4. Testing Compatibility for the Peer-to-Peer System

In this section, we give a sample test case generated for a system consisting of 3 peers and a broker. We have already demonstrated in Table I that model checking the whole system runs out of system resources and is not feasible for 3 or more peers. The test case in Fig. 10 is generated from the composition of the behavioral interfaces of peer and broker (cf. Fig. 3) considering three peers. Execution of this test case takes less than a second.

As mentioned earlier, the server side behavioral interface of a peer allows only one request at a time. This means that two `SReq` messages may not be sent to the same peer before it has replied to the first one. The scenario captured in this test case is designed to check this property for peer number 2 (in Step 5).

This test case starts with registering the servers at the broker followed by requests from clients to the broker. Then the broker replies to these requests by sending via the `server` variable. Since `server` is defined as a `meta` variable, the test case uses a temporary variable `ps` in order to pass on this value to the clients. If a behavioral interface specification requires special conditions on the values of this parameter, the test case would also check these conditions. Finally, if two `SReq` messages are sent to peer number 2, the test fails, i.e., there is a counter-example to compatibility. The test passes if server 2 replies to the first request.

Given the nondeterministic selection of a server in the `CReq` method (cf. Fig. 4), the test case in Fig. 10 will fail. The reason is that the broker may assign the same server to multiple clients which may independently send a request to this server. One simple solution to this

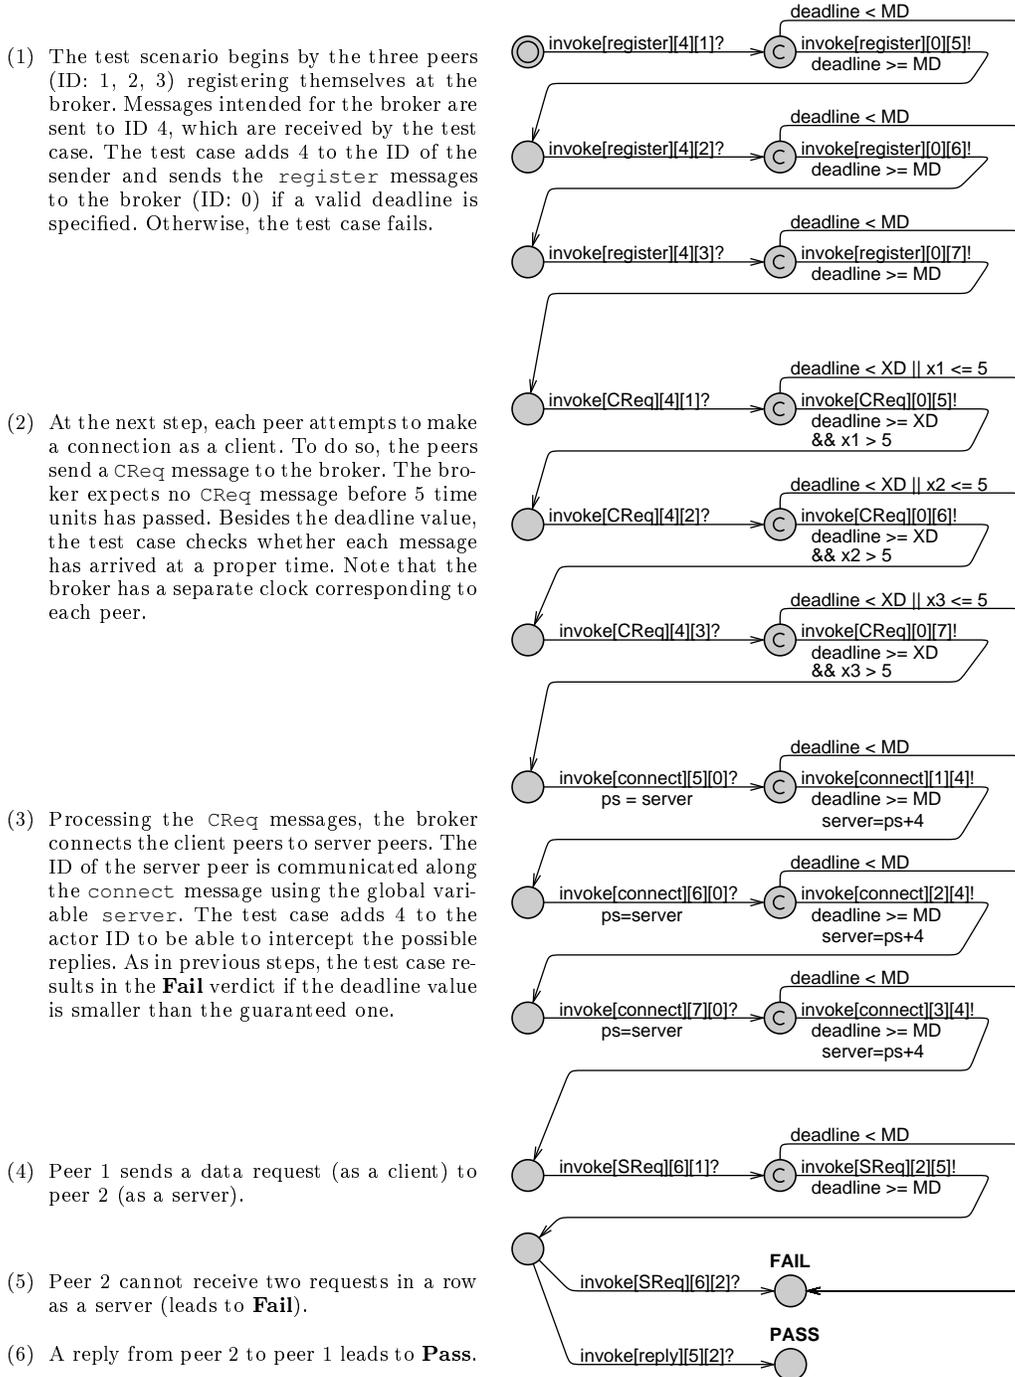


Fig. 10. In this test case, the broker actor is assigned 0 while the peer actors are assigned 1, 2 and 3. In order to intercept the messages between actors, the test automaton represents each actor by adding 4 to its ID (when binding the known actors). For the sake of simplicity, this test case does not include all possible violations of refinement that lead to the **Fail** verdict.

incompatibility would be using a round-robin assignment of the servers by the broker to the incoming requests. With this strategy, the test case in Fig. 10 does not lead to the **Fail** verdict anymore.

## 6. RELATED WORK

There has been lots of work on scheduling in real-time systems [Buttazzo 2011]. The main aspect of our work is that we address schedulability at a *modeling* level as in [Fersman et al. 2007; Altisen et al. 2002; Garcia et al. 2000; Nigro and Pupo 2001], whereas [Closse et al. 2001; Kloukinas and Yovine 2003] are applied to *programming* languages. This results in a major methodological difference. In the latter case, analysis is performed only after the software has been developed: a given application is augmented with real-time requirements (like deadlines) and automata are derived from code. This approach can be useful specifically for legacy software. In contrast, we use automata for platform-independent modeling of actors and their behavioral interfaces at the design stage. One can thus boost the application performance by fine-tuning the scheduling in the early steps of development and at a much lower cost. With this in mind, we compare our schedulability and testing methods with some most relevant works, focusing on its different aspects.

*Task Specifications.* Schedulability analysis results depend on the assumptions on task generation. Many approaches are based on simple task generation patterns like periodic tasks and rate-monotonic analysis [Shin and Lee 2008; Garcia et al. 2000]. Although useful in many cases, such techniques are too coarse in many distributed systems and produce pessimistic outcomes. In our work, behavioral interfaces specify how tasks may be generated in an actor. Being based on Task automata [Fersman et al. 2007], we can describe non-uniformly recurring tasks, inducing much more accurate analysis. We extend also the decidability results by Fersman et al. [Fersman et al. 2007] and show that our analysis (based on non-preemptive scheduling) can be reduced to checking for reachability in timed automata, and is therefore decidable [Henzinger et al. 1994]. Although [Closse et al. 2001; Kloukinas and Yovine 2003] also use automata, they do not discuss decidability.

The main difference between our work and task automata is that in our framework tasks are specified as timed automata. Therefore tasks can trigger other (internal) tasks during execution, which may inherit the (remaining) deadline of the task generating them (called delegation). In task automata, a task is completely abstracted away into an execution time, and generation of all tasks is captured in the task automaton. In our approach, internal tasks cannot be captured in the behavioral interfaces, because their arrival depends on the scheduling of the parent tasks, which in turn depends on the selected scheduling strategy. Our approach is therefore strictly more expressive than task automata.

*Modularity.* Schedulability has usually been analyzed for a whole system running on a single processor, whether at modeling [Fersman et al. 2007; Altisen et al. 2002] or programming level [Closse et al. 2001; Kloukinas and Yovine 2003]. We address distributed systems where each actor has a dedicated processor and scheduling policy. In our approach, behavioral interfaces are key to modularity. They model the most general message arrival pattern for actors. They can be viewed as a contract as in ‘design by contract’ [Meyer 1992] or as a most general assumption in modular model checking [Kupferman et al. 2001] (based on assume-guarantee reasoning); schedulability is guaranteed if the real use of the actor satisfies this assumption in the behavioral interface.

The approach in [Nigro and Pupo 2001] is modular in the sense that the untimed specification of the actors, and the timing constraints (specified separately) can be reused. However, they still analyze a complete system, rather than individual actors. Furthermore, a deadline in their framework includes only the time until an event is received. Hence, their approach cannot address complications like delegation of a task to subtasks. Another related work is TAXYS [Closse et al. 2001], where an abstract model of the environment is used for

schedulability analysis. However, it is used to analyze a complete program and is not used modularly.

*Interface Design.* Our behavioral interfaces are similar to Timed Interface Automata [de Alfaro et al. 2002], but the notion of compatibility is different. Alfaro et al. [de Alfaro et al. 2002] take an optimistic approach in which two interfaces are compatible if there is a possible way for them to work properly. This leads to a simpler theory but to implement these interfaces, one needs to adhere to these possibilities to end up with a working system. David et al. [David et al. 2010] suggest to make specifications input-enabled by adding an Error state and directing every undesired behavior to that state. They define two specifications to be compatible if their composition does not reach the Error state. This is unfortunately too restrictive for high-level specifications; abstract behavioral interfaces easily fall into spurious incompatibilities whereas their implementations may still work together. Our approach bridges the gap between these two methods by considering the actual implementation of actors. We check whether the implementations at hand, when composed, indeed follow the behavior that makes their interfaces compatible (w.r.t. the optimistic approach of [de Alfaro et al. 2002]). Finally, timed actor interfaces in [Geilen et al. 2011] are defined to accept *early* actions. This is an orthogonal feature and can be combined with our notion of behavioral interfaces if desired.

Analyzing the composition of the concurrent objects is subject to state space explosion because of their asynchronous nature and all their queues. We discussed in [Jaghoori 2011] the necessary conditions for compositional model checking of refinement. However it is not sufficient to prove refinement in all cases. We proposed in this paper a sound and complete testing technique for compatibility, based on finding counter-examples to refinement, as positioned below.

*Testing Real-Time Systems.* In real-time systems, conformance (or refinement) is tested in terms of allowed actions as well as of right timings. Different conformance relations have been investigated: timed bisimulation [Cardell-Oliver and Glover 1998], may and must preorders [Nielsen and Skou 2001], timed trace inclusion [Khoumsi et al. 2004; Cattani and Kwiatkowska 2005; Hessel et al. 2008], and timed extension of Tretmans' conformance relation *ioco* [Tretmans 1996; Krichen and Tripakis 2009; Schmaltz and Tretmans 2008]. A naive approach to testing refinement is to take a trace from the concrete system and check whether it exists also in the abstract specification. This approach is impractical because the concrete system is too complicated. Therefore, as is usual in the literature, e.g., in *ioco* testing, we generate test cases from the abstract specification. A main difference is that these techniques are for testing an open system, i.e., by feeding inputs and observing the outputs. However, we deal with a closed system. In this respect, we take advantage of the fact that the system under test is also a model in our case, so we can use a tool like UPPAAL, and drive the system execution by synchronizing on the actions of the test case. This greatly restricts the system behavior, although still nondeterministic. The restricted system behavior is now amenable to model checking in order to address the remaining nondeterminism during test submission.

A similar approach of using testing techniques to avoid state space explosion in the analysis of real-time system models has been followed by Clarke and Lee [Clarke and Lee 1997], in the setting of a real-time process algebraic formalism called ACSR (the Algebra of Communicating Shared Resources). They focus on testing timing constraints of real-time systems, deriving time-efficient test cases from a graphical representation of those constraints and defining time domains coverage criteria.

The soundness of our testing method depends on the determinism of the behavioral interfaces. The problem of determinizability of arbitrary timed automata is undecidable [Finkel 2006; Tripakis 2006], so in this paper we require the behavioral interfaces to be given as deterministic automata. To relax this requirement, one can consider the class

of determinizable timed automata as in [Khoumsi et al. 2004], or use digital test cases [Krichen and Tripakis 2009], where time is discrete to answer the implementability of test cases. Alternatively, automated over-approximation techniques as in [Bertrand et al. 2011] can be employed.

As a final note, timelocks (and sometimes also deadlocks) are generally considered errors in a specification. When testing, the specification and the implementation are usually assumed to be non-blocking, meaning that they will never block time in any environment. However, since our specification is obtained by synchronous product of behavioral interfaces, such cases can happen, it makes sense to allow the same “errors” in the system as in its specification: Our conformance relation then takes into account the presence of deadlocks and timelocks, and allows them in the system whenever they exist in the specification.

## 7. CONCLUSIONS

The main contribution of this work is the integration of the abstract formalism of timed automata into a high-level object based modeling paradigm. On the one hand, the abstraction level of automata theory enables us to provide powerful analysis techniques and specifically less pessimistic schedulability analysis compared to traditional approaches. On the other hand, we augment the successful actor-based approach to object-orientation (as in Scala and Erlang) with application-level scheduling, as also motivated in resource-aware programming techniques.

We presented a complete framework for modular schedulability analysis of distributed systems. Schedulability of each actor is analyzed individually with respect to its behavioral interface. This is made feasible by putting a finite bound on the task queue such that the schedulability results hold for any queue length. We can then test a system of communicating objects to make sure objects are used as expected. This compatibility further implies the schedulability of the whole system. In this paper, we specifically gave a detailed account of a novel counter-example oriented technique for testing refinement as the basis for compatibility check.

As future work, we are planning to integrate this high-level analysis framework into our implementation of application-level scheduling on top of Java. The integrated tool suite will span a complete software development cycle, and will be a basis for developing safety critical real-time distributed and embedded systems. The main advantage of such a tool suite is that the designer/programmer will be in control of scheduling in the whole development cycle, which is in turn the key to efficiency in the software running on future multi-core and cloud infrastructures.

A specific and interesting case where we can apply our approach is in modeling TinyOS [Hill et al. 2000; Cheong 2007]. TinyOS is an actor-based open-source runtime environment designed for sensor network and has a large user base of over 500 research groups and companies [Cheong 2007]. The event-driven execution model of TinyOS enables fine-grained power management yet allows the scheduling flexibility made necessary by the unpredictable nature of wireless communication and physical world interfaces. Modeling a TinyOS instance as an actor, we can define its own scheduling policy and hence the designer is able to introduce and analyze different policies in scheduling.

## REFERENCES

- ACHARYA, A., RANGANATHAN, M., AND SALTZ, J. H. 1996. Sumatra: A language for resource-aware mobile programs. In *Mobile Object Systems*. 111–130.
- AGHA, G. 1990. The structure and semantics of actor languages. In *Proc. the REX Workshop*. 1–59.
- AGHA, G., MASON, I., SMITH, S., AND TALCOTT, C. 1997. A foundation for actor computation. *Journal of Functional Programming* 7, 1–72.
- ALTISEN, K., GÖSSLER, G., AND SIFAKIS, J. 2002. Scheduler modeling based on the controller synthesis paradigm. *Real-Time Systems* 23, 1-2, 55–84.

- ALUR, R. AND DILL, D. L. 1994. A theory of timed automata. *Theoret. Comp. Sci.* 126, 2, 183–235.
- ALUR, R. AND WEISS, G. 2008. Rtcomposer: a framework for real-time components with scheduling interfaces. In *Proc. Embedded software (EMSOFT'08)*, L. de Alfaro and J. Palsberg, Eds. ACM, 159–168.
- ARMSTRONG, J. 2010. Erlang. *Communications of ACM* 53, 9, 68–75.
- BERMAN, F. AND WOLSKI, R. 1996. Scheduling from the perspective of the application. In *Proc. High Performance Distributed Computing (HPDC'96)*. IEEE Computer Society, 100–111.
- BERTRAND, N., STAINER, A., JÉRON, T., AND KRICHEN, M. 2011. A game approach to determinize timed automata. In *Proc. Foundations of software science and computational structures. FOSACS'11/ETAPS'11*. Springer-Verlag, 245–259.
- BUTTAZZO, G. 2011. *Hard Real-Time Computing Systems* Third Ed. Springer.
- CARDELL-OLIVER, R. AND GLOVER, T. 1998. A practical and complete algorithm for testing real-time systems. In *Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'98)*. Lecture Notes in Computer Science Series, vol. 1486. 251–261.
- CATTANI, S. AND KWIATKOWSKA, M. Z. 2005. A refinement-based process algebra for timed automata. *Formal Asp. Comput.* 17, 2, 138–159.
- CHANG, P.-H. AND AGHA, G. 2007a. Supporting reconfigurable object distribution for customized web applications. In *The 22nd Annual ACM Symposium on Applied Computing (SAC)*. 1286–1292.
- CHANG, P.-H. AND AGHA, G. 2007b. Towards context-aware web applications. In *7th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS)*. 239–252.
- CHEONG, E. 2007. Actor-oriented programming for wireless sensor networks. Ph.D. thesis, Electrical Engineering and Computer Sciences University of California at Berkeley.
- CHEONG, E., LEE, E. A., AND ZHAO, Y. 2005. Viptos: a graphical development and simulation environment for tinyos-based wireless sensor networks. In *Proc. Embedded net. sensor sys., SenSys 2005*. 302–302.
- CLARKE, D. AND LEE, I. 1997. Automatic test generation for the analysis of a real-time system: Case study. In *IEEE Real Time Technology and Applications Symposium*. 112–124.
- CLOSSE, E., POIZE, M., PULOU, J., SIFAKIS, J., VENTER, P., WEIL, D., AND YOVINE, S. 2001. TAXYS: A tool for the development and verification of real-time embedded systems. In *Proc. Computer Aided Verification*, G. Berry, H. Comon, and A. Finkel, Eds. LNCS Series, vol. 2102. Springer, 391–395.
- COURCOUBETIS, C. AND YANNAKAKIS, M. 1992. Minimum and maximum delay problems in real-time systems. *Formal Methods in System Design* 1, 4, 385–415.
- DAVID, A., LARSEN, K. G., LEGAY, A., NYMAN, U., AND WASOWSKI, A. 2010. Timed I/O automata: a complete specification theory for real-time systems. In *Proc. Hybrid Systems: Computation and Control (HSCC'10)*. ACM, 91–100.
- DE ALFARO, L., HENZINGER, T. A., AND STOELINGA, M. 2002. Timed interfaces. In *Proc. Embedded Software (EMSOFT)*. LNCS Series, vol. 2491. 108–122.
- FERSMAN, E., KRICAL, P., PETERSSON, P., AND YI, W. 2007. Task automata: Schedulability, decidability and undecidability. *Information and Computation* 205, 8, 1149–1172.
- FINKEL, O. 2006. Undecidable problems about timed automata. In *Proc. Formal Modeling and Analysis of Timed Systems (FORMATS'06)*, E. Asarin and P. Bouyer, Eds. LNCS Series, vol. 4202. Springer, 187–199.
- GARCIA, J. J. G., GUTIERREZ, J. C. P., AND HARBOUR, M. G. 2000. Schedulability analysis of distributed hard real-time systems with multiple-event synchronization. In *Proc. 12th Euromicro Conference on Real-Time Systems*. IEEE, 15–24.
- GEILEN, M., TRIPAKIS, S., AND WIGGERS, M. 2011. The earlier the better: a theory of timed actor interfaces. In *Proc. Hybrid Systems: Computation and Control (HSCC'11)*, M. Caccamo, E. Frazzoli, and R. Grosu, Eds. ACM, 23–32.
- GRABE, I., JAGHOORI, M. M., KLEIN, J., KLÜPPELHOLZ, S., STAM, A., BAIER, C., BLECHMANN, T., AICHERNIG, B. K., DE BOER, F. S., GRIESMAYER, A., JOHNSEN, E. B., KYAS, M., LEISTER, W., SCHLATTE, R., STEFFEN, M., TSCHIRNER, S., LIANG, X., AND YI, W. 2009. The credo methodology - (extended version). In *Proc. 8th Formal Methods for Components and Objects (FMCO'09)*, F. S. de Boer, M. M. Bonsangue, S. Hallerstede, and M. Leuschel, Eds. LNCS Series, vol. 6286. 41–69.
- HALLER, P. AND ODERSKY, M. 2009. Scala actors: Unifying thread-based and event-based programming. *Theoretical Computer Science* 410, 2-3, 202–220.
- HENZINGER, T. A., NICOLLIN, X., SIFAKIS, J., AND YOVINE, S. 1994. Symbolic model checking for real-time systems. *Inf. Comput.* 111, 2, 193–244.
- HESSEL, A., LARSEN, K. G., MIKUCIONIS, M., NIELSEN, B., PETERSSON, P., AND SKOU, A. 2008. Testing real-time systems using uppaal. In *Formal Methods and Testing*. LNCS Series, vol. 4949. 77–117.

- HEWITT, C. 1971. Procedural embedding of knowledge in planner. In *Proc. the 2nd International Joint Conference on Artificial Intelligence*. 167–184.
- HEWITT, C. 2007. What is commitment? physical, organizational, and social (revised). In *Proc. Coordination, Organizations, Institutions, and Norms in Agent Systems II*. LNCS Series. Springer, 293–307.
- HILL, J., SZEWCZYK, R., WOO, A., HOLLAR, S., CULLER, D., AND PISTER, K. 2000. System architecture directions for networked sensors. In *Proc. Arch. Support for Prog. Lang. and Operating Sys.* 93–104.
- JAGHOORI, M. M. 2011. Composing real-time concurrent objects - refinement, compatibility and schedulability. In *Proc. Fundamentals of Soft. Eng. (FSEN'11)*. LNCS Series, vol. 7141. Springer, 96–111.
- JAGHOORI, M. M., DE BOER, F. S., CHOTHIA, T., AND SIRJANI, M. 2009. Schedulability of asynchronous real-time concurrent objects. *J. Logic and Alg. Prog.* 78, 5, 402 – 416.
- JAGHOORI, M. M., LONGUET, D., DE BOER, F. S., AND CHOTHIA, T. 2008. Schedulability and compatibility of real time asynchronous objects. In *Proc. Real Time Systems Symposium*. IEEE CS, 70–79.
- JARD, C., JÉRON, T., AND MOREL, P. 2000. Verification of test suites. In *International Conference on Testing Communicating Systems (TestCom 2000)*. 3–18.
- JOHNSON, E. B. AND OWE, O. 2007. An asynchronous communication model for distributed concurrent objects. *Software and Systems Modeling* 6, 1, 35–58.
- KARMANI, R. K., SHALI, A., AND AGHA, G. 2009. Actor frameworks for the jvm platform: a comparative analysis. In *Proc. Principles and Practice of Prog. in Java (PPPJ'09)*. ACM, 11–20.
- KHOUMSI, A., JÉRON, T., AND MARCHAND, H. 2004. Test cases generation for nondeterministic real-time systems. In *Formal Approaches to Software Testing (FATES'03)*. LNCS Series, vol. 2931. 131–146.
- KLOUKINAS, C. AND YOVINE, S. 2003. Synthesis of safe, QoS extendible, application specific schedulers for heterogeneous real-time systems. In *Proc. Euromicro Conf. on Real-Time Sys.* IEEE CS, 287–294.
- KRICHEN, M. AND TRIPAKIS, S. 2004. Black-box conformance testing for real-time systems. In *Model Checking Software, 11th International SPIN Workshop*. LNCS Series, vol. 2989. 109–126.
- KRICHEN, M. AND TRIPAKIS, S. 2009. Conformance testing for real-time systems. *Formal Methods in System Design* 34, 3, 238–304.
- KUPFERMAN, O., VARDI, M. Y., AND WOLPER, P. 2001. Module checking. *Information and Computation* 164, 2, 322–344.
- LARSEN, K. G., PETTERSSON, P., AND YI, W. 1997. UPPAAL in a nutshell. *STTT* 1, 1-2, 134–152.
- LEE, E. A., LIU, X., AND NEUENDORFFER, S. 2009. Classes and inheritance in actor-oriented design. *ACM Transactions in Embedded Computing Systems* 8, 4.
- LEE, E. A., NEUENDORFFER, S., AND WIRTHLIN, M. J. 2003. Actor-oriented design of embedded hardware and software systems. *Journal of Circuits, Systems, and Computers* 12, 3, 231–260.
- MACKENZIE, K. AND WOLVERSON, N. 2003. Camelot and grail: resource-aware functional programming for the jvm. In *Trends in Functional Programming*. 29–46.
- MEYER, B. 1992. *Eiffel: The language*. Prentice-Hall. (first printing: 1991).
- MOREAU, L. AND QUEINNEC, C. 2005. Resource aware programming. *ACM Trans. Program. Lang. Syst.* 27, 3, 441–476.
- NIELSEN, B. AND SKOU, A. 2001. Automated test generation from timed automata. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS'01)*. LNCS Series, vol. 2031. 343–357.
- NIGRO, L. AND PUPO, F. 2001. Schedulability analysis of real time actor systems using coloured petri nets. In *Proc. Concurrent Object-Oriented Prog. and Petri Nets*. LNCS Series, vol. 2001. Springer, 493–513.
- NOBAKHT, B., DE BOER, F. S., JAGHOORI, M. M., AND SCHLATTE, R. 2012. Programming and deployment of active objects with application-level scheduling. In *Proc. ACM Symposium on Applied Computing (SAC'12)*. ACM. To appear.
- REED, G. AND ROSCOE, A. 1999. The timed failures and stability model for csp. *Theoretical Computer Science* 211, 1–2, 85 – 127.
- SCHMALTZ, J. AND TRETMANS, J. 2008. On conformance testing for timed systems. In *Formal Modeling and Analysis of Timed Systems*. LNCS Series, vol. 5215. Springer, 250–264.
- SHIN, I. AND LEE, I. 2008. Compositional real-time scheduling framework with periodic model. *ACM Trans. Embed. Comput. Syst.* 7, 30:1–30:39.
- SIMONS, D. P. L. AND STOELINGA, M. 2001. Mechanical verification of the IEEE 1394a root contention protocol using Uppaal2k. *STTT* 3, 4, 469–485.
- SIRJANI, M., MOVAGHAR, A., SHALI, A., AND DE BOER, F. S. 2004. Modeling and verification of reactive systems using Rebeca. *Fundamenta Informaticae* 63, 4, 385–410.
- TRETMANS, J. 1996. Test generation with inputs, outputs and repetitive quiescence. *Software - Concepts and Tools* 17, 3, 103–120.

TRIPAKIS, S. 2006. Folk theorems on the determinization and minimization of timed automata. *Inf. Process. Lett.* 99, 222–226.

## A. SEMANTICS OF ACTOR SYSTEMS

Consider a set of actors  $O_1, \dots, O_n$  implementing the behavioral interfaces  $B_1, \dots, B_n$ , respectively. This set of actors comprises a closed system if for every actor  $O_j$ , we have  $Act_j^O \subseteq (M_{O_j} \cup \bigcup_{1 \leq i \leq n} M_{B_i})$  where  $Act_j^O$  is the set of calls made by methods of  $O_j$ .

Before defining the semantics of such a system, we repeat the definition of a scheduler from [Jaghoori et al. 2009].

*Definition A.1 (Scheduler Function).* A scheduler function ‘ $sched(q, m(d))$ ’ is a function which given a queue  $q$  with clocks  $C_q$  and a method name  $m$  with a deadline  $d$ , returns a set of triples  $\{(G, c, q')\}$ , where

- $G$  is a guard on clocks in  $C_q$  (possibly based on  $d$ );
- $c \in C_q$  is a clock not used (i.e., not assigned to any tasks) in  $q$ ; and,
- $q'$  is a queue with the clocks  $C_q$  and represents the queue  $q$  after inserting  $m(d, c)$  in a particular position as implied by the guard  $G$ .

An overloading of the scheduler function is defined as  $sched(q, m(d, c))$  such that it inserts a task into the queue using a given clock  $c$ . By reusing the deadline and the clock already assigned to a task in the queue, we can model inheriting the deadline. In the case of delegation, the clock assigned to the currently running task is reused.

A scheduler function is *preemptive* if it can place the new task in the first position. Recall that the first task in the queue is the task that is currently running. In this paper we only consider non-preemptive schedulers.

The semantics of a system is defined by a timed automaton, called the system automaton. The *system automaton* for a given system with method names  $\mathcal{M}_S = \bigcup_{1 \leq i \leq n} M_{O_i}$  and a scheduler function  $sched$  is a timed automaton  $S = (L_S, l_S, E_S, I_S)$  over the alphabet  $Act_S = \mathcal{M}_S$  and the clocks  $C_S$ :

- The set of clocks  $C_S$  is the union of all sets of clocks for the method automata plus the queue clocks of each object.
- The locations of the system are the product of each of the locations of actors together with a queue, i.e.,  $\{(l_1, q_1), (l_2, q_2), \dots, (l_n, q_n)\}$ .
- The initial location  $l_S$  is:

$$\{ (start(A_1), [m_1(d_1, c_1)]), \dots, (start(A_n), [m_n(d_n, c_n)]) \}$$

- where  $m_i: A_i$  is the ‘initial’ method of the actor  $O_i$ ,  $d_i$  is the deadline for this initial method and  $c_i$  is one of the object’s queue clocks.
- The edges  $E_S$  are defined with the rules in Fig. 11 ( $M_l$  is the set of methods provided by the object that is in location  $l$ ). We write  $\xrightarrow[g; r]{m}_S$  for an edge of  $S$  with action  $m$ , guard  $g$  and update  $r$ .
- The invariant of a location is defined as the conjunction of the location invariants of all currently executing object locations.

In Fig. 11, function  $start(m)$  returns the initial location of the automaton for method  $m$ . Locations in method automata with no outgoing transitions are called *final*. The first three rules in this figure take care of enqueueing sent messages. The first two are for self calls and therefore the queue of the same object is used. In case of delegation, the clock of the current task is reused and thus the deadline is inherited. Whenever the execution of the current task finishes, the context switch rule makes sure the next method in the queue is executed,

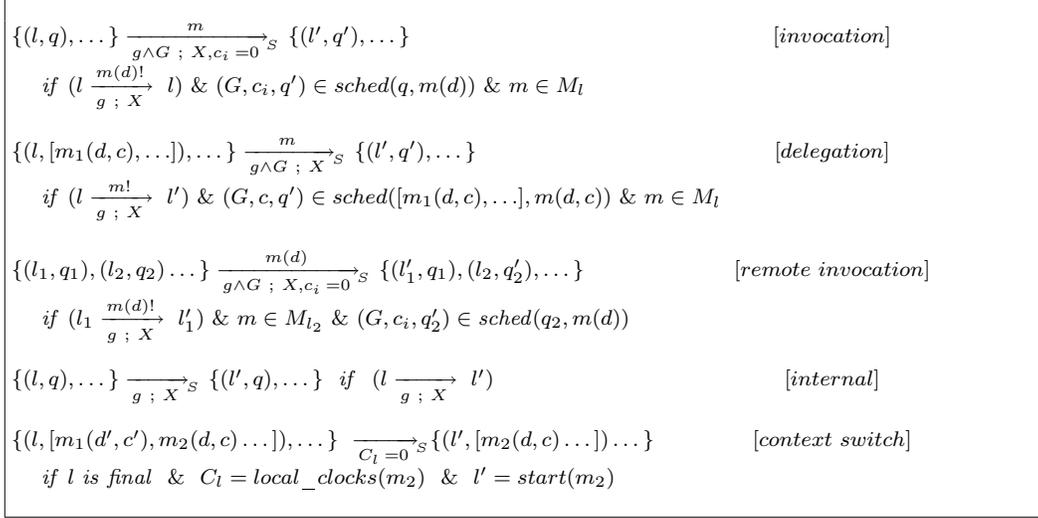


Fig. 11. Reductions for a System

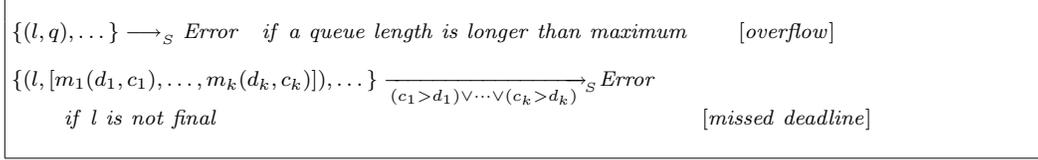


Fig. 12. Missing a deadline leads to an Error

if there is any. If in a location  $\{(l, q), \dots\}$ ,  $l$  is final and  $q$  has more than one element, the location is marked urgent. This forces context switch to happen as soon as it is possible.

In order to enable model checking as described in Section 4, we need to put a maximum on the queue length and ensure the rules in Fig. 11 are used only if the queue length is less than the maximum. Going beyond the maximum queue length for any actor or missing a deadline by a message in a queue must lead to an error state. These extra rules are depicted in Fig. 12.

## B. MODELING SCHEDULERS IN UPPAAL

A scheduler function, as described in the previous section, can be implemented as a scheduler automaton. This automaton also contains a queue. Fig. 13 shows the general structure of a scheduler automaton. This general picture does not specify any specific scheduling strategy. The scheduler automata applies the scheduling strategy at dispatch time instead of insertion time, but the resulting behavior is the same. The reason is to enable using deadlines in the strategy. As explained in the Section 3, the deadline value cannot be used (in the guard) on the same transition where a message is received.

*Queue.* Tasks in the queue are modeled using the following arrays:  $q$  holds the message names,  $d$  holds their initial deadline values and  $clk$  consists of clocks that keep track of the time a task has been in the queue. The sender of every message is stored in the  $s$  array. If messages can have parameters,  $p$  arrays will be added for each parameter. We assume a maximum length of  $MAX$  for these arrays. As described in Section 4, we can find such a maximum for queues of schedulable objects. The array  $ca$  shows the clock assigned to each message (task), such that ' $d[ca[i]] - clk[ca[i]]$ ' represents the remaining deadline of  $q[i]$

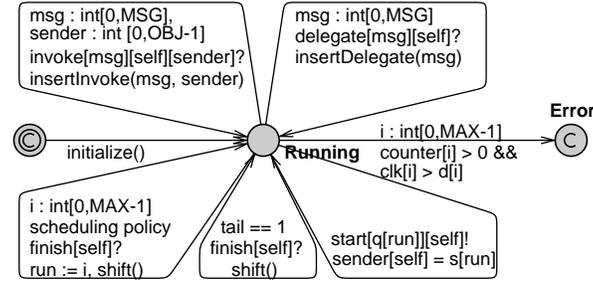


Fig. 13. A general scheduler automaton (repeated from the main text)

```

1 void initialize(){
2     q[0] = op_init;
3     s[0] = self;
4     //p[0] remains 0
5     d[0] = INIT_DEADLINE;
6     ca[0] = 0;
7     // clock clk[0] is assigned
8     counter[0] = 1;
9     tail = 1;
10 }
11 void shift(int a) {
12     counter[ca[a]] --;
13     if (counter[ca[a]] == 0)
14         // if the clock is no longer used
15         d[ca[a]] = MD;
16     tail --;
17     while (a < tail && a < MAX-1) {
18         q[a] = q[a+1];
19         ca[a] = ca[a+1];
20         a++;
21     }
22     q[a] = 0;
23     ca[a] = 0;
24 }
25 void insertInvoke(int m, int snd){
26     int c, i;
27     c = MAX;
28     for (i = 0; c == MAX; i++) {
29         if (counter[i][s]==0) c = i;
30     }
31     q[tail] = m;
32     s[tail] = snd;
33     // store parameters if any, e.g.
34     // p[tail] = server; in Peer
35     ca[tail] = c;
36     clk[c] = 0;
37     d[c] = deadline;
38     counter[c] = 1;
39     tail++;
40 }
41 void insertDelegate(int msg){
42     q[tail] = msg;
43     s[tail] = s[run];
44     // pass on parameters if any
45     ca[tail] = ca[run];
46     counter[ca[tail]] ++;
47     tail ++;
48 }

```

Fig. 14. Inserting a message into the queue using `invoke` and `delegate` mechanisms

at any time. `counter[i]` holds the number of tasks using clock `clk[i]`. A clock is free if its counter is zero. When delegation is used, the counter becomes greater than one.

*Initialization.* The initialization of a queue takes place in the `initialize` function. This transition is taken before any method in any actor is started, because its start location is committed. The function shown in Fig. 14 puts the initialization method `op_init` in the queue and assigns the first free clock to it.

*Input-enabledness.* A scheduler for a class  $R$  should allow receiving any message in  $M_R$  at any time. In Fig. 13, there is an edge (top-left in the picture) that allows receiving a message on the `invoke` channel (from any sender). To allow any message and sender, ‘select’

expressions are used. The expression `msg : int[0,MSG]` nondeterministically selects a value between 0 and `MSG` for `msg`. This is equivalent to adding a transition for each value of `msg`. Similarly, any sender (`sender : int[0,OBJ-1]`) can be selected. This message is put at the tail of the queue (`q[tail] = msg`), and a free clock (`counter[c] == 0`) is assigned to it (`ca[tail] = c`), and the deadline value is recorded (`d[c] = deadline`); this is handled in the function `insertInvoke` shown in Fig. 14. The synchronization between this transition and the method automata corresponds to the *invocation* rules in Fig. 11.

A similar transition accepts messages on the `delegate` channel (top-right in the picture). In this case, the clock already assigned to the currently running task (parent task) is assigned to the internal task (`ca[tail] = ca[run]`); this is handled in the function `insertDelegate` shown in Fig. 14. In a delegated task, no sender is specified (it is always `self`). The variable `run` shows the index of the currently running task in the queue (which is not necessarily the first task). This handles the rule *delegation* in Fig. 11.

*Error.* The scheduler automaton moves to the `ERROR` state if a deadline is missed (`clk[i] > d[i]`). The guard `counter[i] > 0` checks whether the corresponding clock is currently in use, i.e., assigned to a message in the queue. Furthermore, to make sure no queue overflow occurs, the property to check should include  $tail \leq MAX$ .

*Scheduling Strategy.* When a message is added to an empty queue, the corresponding method is immediately started. When a method is finished (synchronizing on `finish` channel), it is taken out of the queue (by `shift()` given in Fig. 14). If the currently running method is the last in the queue, nothing needs to be selected (i.e., if `tail == 1` we only need to `shift`). Otherwise, the next method to be executed should be chosen based on a specific scheduling strategy (by assigning the right value to `run`). For a *concrete* scheduler, the guard and update of `run` should be well defined. If `run` is always assigned 0 during context switch, the automaton serves as a First Come First Served (FCFS) scheduler. In an FCFS scheduler, the two transitions on `finish` channel can be combined.

A Fixed Priority Scheduler (FPS) can be implemented by associating a constant priority value to each method/task type. Suppose the array `p` represents the static priority of all methods, such that for a message `q[i]` in the queue, its priority can be obtained by `p[q[i]]`. We can then formulate FPS strategy using the guard:

```
i < tail && i != run &&
forall (m : int[0,MAX-1])
  (m == run) || (p[q[i]] >= p[q[m]])
```

This formula selects `i` such that `q[i]` is not an empty queue cell (`i < tail`) or the currently finished method (`run`), and `p[q[i]]` is the highest priority. If there are multiple tasks with the same priority, the following can be added to the above guard to ensure the first task is selected:

```
forall (m : int[0,MAX-1])
  m == running || m <= i || m >= tail || x[ca[i]]-x[ca[m]] >= d[ca[i]]-d[ca[m]]
```

An Earliest Deadline First (EDF) scheduler always selects the task with the smallest remaining deadline. This is an example of dynamic priority scheduling because the remaining deadline of a task gets smaller as time passes. The remaining deadline of message `i` is given by `d[i] - clk[i]`. This can be encoded using a guard like:

```
i < tail && i != run &&
forall (m : int[0,MAX-1])
  (m == run) || (clk[ca[i]] - clk[ca[m]] >= d[ca[i]] - d[ca[m]])
```

and  $i$  will show the task with the smallest remaining deadline. Notice that  $clk[a] - clk[m] \geq d[a] - d[m]$  is equivalent to  $d[m] - clk[m] \geq d[a] - clk[a]$ . The rest ensures that an empty queue cell ( $i < tail$ ) or the currently finished method ( $run$ ) is not selected.

After the next method to execute is selected, context-switch happens by starting the selected method. Having defined `start` as an urgent channel, the next method is immediately scheduled (if queue is not empty) by taking the bottom-right transition in Fig. 13.