BRICS

Basic Research in Computer Science

Model Checking via Reachability Testing for Timed Automata

Luca Aceto
Augusto Burgueño
Kim G. Larsen

BRICS Report Series **RS-97-29**

ISSN 0909-0878 **November 1997**

Copyright © 1997, BRICS, Department of Computer Science University of Aarhus. All rights reserved.

Reproduction of all or part of this work is permitted for educational or research use on condition that this copyright notice is included in any copy.

See back inner page for a list of recent BRICS Report Series publications. Copies may be obtained by contacting:

BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK–8000 Aarhus C
Denmark

Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk

BRICS publications are in general accessible through the World Wide Web and anonymous FTP through these URLs:

`http://www.brics.dk`
`ftp://ftp.brics.dk`
This document in subdirectory `RS/97/29/`

Model Checking via Reachability Testing for Timed Automata

Luca Aceto[1] *, Augusto Burgueño[2] ** and Kim G. Larsen[1]

[1] BRICS***, Department of Computer Science, Aalborg University,
Fredrik Bajers Vej 7-E, DK-9220 Aalborg Ø, Denmark.
Email: {luca,kgl}@cs.auc.dk, Fax: +45 98 15 98 89
[2] ONERA-CERT, Département d'Informatique,
2 av. E. Belin, BP4025, 31055 Toulouse Cedex 4, France.
Email: a.burgueno@acm.org, Fax: +33 5 62 25 25 93

Abstract. In this paper we develop an approach to model-checking for timed automata via reachability testing. As our specification formalism, we consider a dense-time logic with clocks. This logic may be used to express safety and bounded liveness properties of real-time systems. We show how to automatically synthesize, for every logical formula φ, a so-called *test automaton* T_φ in such a way that checking whether a system S satisfies the property φ can be reduced to a reachability question over the system obtained by making T_φ interact with S.
The testable logic we consider is both of practical and theoretical interest. On the practical side, we have used the logic, and the associated approach to model-checking via reachability testing it supports, in the specification and verification in UPPAAL of a collision avoidance protocol. On the theoretical side, we show that the logic is powerful enough to permit the definition of *characteristic properties*, with respect to a timed version of the ready simulation preorder, for nodes of deterministic, τ-free timed automata. This allows one to compute behavioural relations via our model-checking technique, therefore effectively reducing the problem of checking the existence of a behavioural relation among states of a timed automaton to a reachability problem.

1 Introduction

Model-checking of real time systems has been extensively studied in the last few years, leading to both important theoretical results, setting the limits of decidability [AD94, HK94, HKPV95], and to the emergence of practical tools as HyTech [HHWT95], Kronos [OY93] and UPPAAL

* Partially supported by the Human Capital and Mobility project EXPRESS.
** Partially supported by Research Grant of the Spanish Ministry of Education and Culture and by BRICS. This work was carried out while the author was visiting Aalborg University.
*** Basic Research in Computer Science.

[BLL+95], which have been successfully applied to the verification of real sized systems [BGK+96, HWT95].

The main motivation for the work presented in this paper stems from our experience with the verification tool UPPAAL. In such a tool, real-time systems are specified as networks of timed automata [AD94], which are then the object of the verification effort. The core of the computational engine of UPPAAL consists of a collection of efficient algorithms that can be used to perform reachability analysis over a model of an actual system. Any other kind of verification problem that the user wants to ask UPPAAL to perform must be encoded as a suitable reachability question. A typical example of such a problem is that of *model checking*. Experience has shown that it is often convenient to describe desired system properties as formulae of some real-time variant of standard modal or temporal logics (see, e.g., [AH94, HNSY94, LLW95]). The model-checking problem then amounts to deciding whether a given system specification has the required property or not.

The way model-checking of properties other than plain reachability ones may currently be carried out in UPPAAL is as follows. Given a property φ to model-check, the user must provide a *test automaton* T_φ for that property. This test automaton must be such that the original system has the property expressed by φ if, and only if, none of the distinguished reject states of T_φ can be reached when the test automaton is made to interact with the system under investigation.

As witnessed by existing applications of this approach to verification by model-checking (cf., e.g., [JLS96]), the construction of a test automaton from a temporal formula or informally specified requirements is a task that, in general, requires a high degree of ingenuity, and is error-prone. It would therefore be useful to automate this process by providing a compilation procedure from formulae in some sufficiently expressive real-time logic into appropriate test automata, and establishing its correctness once and for all. Apart from its practical and theoretical interest, the existence of such a connection between specification logics and automata would also free the average user of a verification tool like UPPAAL from the task of having to generate ad hoc test automata in his/her verifications based on the model-checking approach. We envisage that this will help make the tool usable by a larger community of designers of real-time systems.

1.1 Results

In this paper we develop an approach to model-checking for timed automata via reachability testing. As our specification formalism, we con-

sider a dense-time logic with clocks, which is a fragment of the one presented in [LLW95]. This logic may be used to express safety and bounded liveness properties of real-time systems. We show how to automatically synthesize, for every logical formula φ, a so-called *test automaton* T_φ in such a way that checking whether a system S satisfies the property φ can be reduced to a reachability question over the system obtained by making T_φ interact with S. More precisely, we show that S satisfies property φ iff none of the distinguished reject nodes of the test automaton can be reached in the combined system $S \parallel T_\varphi$ (Thm. 5.2). This result is obtained for a model of timed automata with urgent actions and the interpretation of parallel composition used in UPPAAL.

The logic we consider in this paper only allows for a restricted use of the boolean 'or' operator, and of the diamond modality of Hennessy-Milner logic [HM85]. We argue that these restrictions are necessary to obtain testability of the logic, in the sense outlined above (Propn. 5.4). Indeed, as it will be shown in a companion paper [ABBL97], the logic presented in this study is remarkably close to being completely expressive with respect to reachability properties. In fact, a slight extension of the logic considered here allows us to reduce any reachability property of a composite system $S \parallel T$ to a model-checking problem of S.

Despite the aforementioned restrictions, the testable logic we consider is both of practical and theoretical interest. On the practical side, we have used the logic, and the associated approach to model-checking via reachability testing it supports, in the specification and verification in UPPAAL of a collision avoidance protocol. This protocol was originally analyzed in [JLS96], where rather complex test automata were derived in an ad hoc fashion from informal specifications of the expected behaviour of the protocol. The verification we present here is based on our procedure for the automatic generation of test automata from logical specifications. This has allowed us to turn logical specifications of the expected behaviour of the protocol into automata, whose precise fit with the original properties is guaranteed by construction.

On the theoretical side, we show that the logic is powerful enough to permit the definition of *characteristic properties* [SI94], with respect to a timed version of the ready simulation preorder [LS91, BIM95], for nodes of deterministic, τ-free timed automata (Thm. 8.4). This allows one to compute behavioural relations via our model-checking technique, therefore effectively reducing the problem of checking the existence of a behavioural relation among states of a timed automaton to a reachability problem. As the version of ready simulation we consider preserves the

properties expressible in the logic studied in this paper (Thm. 8.2), our model-checking technique may be used to automatically justify abstraction steps in hierarchical system verifications.

1.2 Related Literature

This study establishes a connection between a logical property language for the specification of safety and bounded liveness properties of real-time systems and the formalism of timed automata. Our emphasis is on the reduction of the model-checking problem for the logic under consideration to an intrinsically automata-theoretic problem, viz. that of checking for the reachability of some distinguished nodes in a timed automaton. The blueprint of this endeavour lies in the automata-theoretic approach to the verification of finite-state reactive systems pioneered by Vardi and Wolper [VW86, VW94, Var96]. In this approach to verification, the intimate relationship between linear time propositional temporal logic [Pnu77] and ω-automata is exploited to yield elegant and efficient algorithms for the analysis of specifications, and for model-checking. The work presented in this paper is not based on a similarly deep mathematical connection between the logic and timed automata (indeed, it is not clear that such a connection exists because, as shown in [AD94], timed Büchi automata are not closed under complementation), but draws inspiration from that beautiful theory. In particular, the avenue of investigation pursued in this study may be traced back to the seminal [VW86].

A characteristic formula construction for timed bisimulation over τ-free timed automata is presented in [LLW95]. Apart from the aforementioned references, the use of characteristic formulae in the computation of behavioural relations is advocated in, e.g., [CS91].

The recent paper [SVD97] presents a generalization of the classical theory of testing for Mealy machines to dense real-time systems. In particular, the authors of *op. cit.* give a test generation algorithm for black-box conformance testing of timed I/O automata. These automata are τ-free, deterministic timed automata that satisfy some additional requirements guaranteeing their testability, and their behaviour can be completely characterized, up to bisimulation equivalence, by a finite collection of finite sequences of actions and delays. The deep connections between this work and the results presented in this paper are still to be explored. We remark, however, that the characteristic formula construction we present for τ-free deterministic timed automata may be seen as a logical formulation of a complete set of experiments, in the sense of [SVD97], that capture a timed version of ready simulation.

1.3 Road-map of the Paper

The paper is organized as follows. We begin by introducing timed automata and timed labelled transition systems (Sect. 2). The notion of test automaton considered in this paper is introduced in Sect. 3, together with the interaction between timed automata and tests. We then proceed to present a real-time logic suitable for expressing safety and bounded liveness properties of real-time systems (Sect. 4). The step from logical properties to test automata is discussed in Sect. 5, and its implementation in UPPAAL in Sect. 6. Section 7 is devoted to a brief description of the specification and verification of a collision avoidance protocol using the theory developed in this paper. The construction of characteristic formulae with respect to timed ready simulation for nodes of deterministic, τ-free automata is the topic of Sect. 8. The paper concludes with a mention of some further results we have obtained on the topic of this paper, and a discussion of interesting subjects for future research (Sect. 9).

2 Preliminaries

We begin by briefly reviewing the timed automaton model proposed by Alur and Dill [AD94].

2.1 Timed Labelled Transition Systems

Let $\mathcal{A}$ be a finite set of actions ranged over by a. We assume that $\mathcal{A}$ comes equipped with a mapping $\overline{\cdot} : \mathcal{A} \to \mathcal{A}$ such that $\overline{\overline{a}} = a$ for every $a \in \mathcal{A}$. We let $\mathcal{A}_\tau$ stand for $\mathcal{A} \cup \{\tau\}$, where τ is a symbol not occurring in $\mathcal{A}$, and use μ to range over it. Following Milner [Mil89], τ will stand for an internal action of a system. Let $\mathbb{N}$ denote the set of natural numbers and $\mathbb{R}_{\geq 0}$ the set of non-negative real numbers. We use $\mathcal{D}$ to denote the set of delay actions $\{\epsilon(d) \mid d \in \mathbb{R}_{\geq 0}\}$, and $\mathcal{L}$ to stand for the union of $\mathcal{A}_\tau$ and $\mathcal{D}$.

Definition 2.1. A *timed labelled transition system* (TLTS) is a structure $\mathcal{T} = \langle S, \mathcal{L}, s^0, \longrightarrow \rangle$ where S is a set of *states*, $s^0 \in S$ is the initial state, and $\longrightarrow \subseteq S \times \mathcal{L} \times S$ is a transition relation satisfying the following properties:

- (TIME DETERMINISM) for every $s, s', s'' \in S$ and $d \in \mathbb{R}_{\geq 0}$, if $s \xrightarrow{\epsilon(d)} s'$ and $s \xrightarrow{\epsilon(d)} s''$, then $s' = s''$;
- (TIME ADDITIVITY) for every $s, s'' \in S$ and $d_1, d_2 \in \mathbb{R}_{\geq 0}$, $s \xrightarrow{\epsilon(d_1+d_2)} s''$ iff $s \xrightarrow{\epsilon(d_1)} s' \xrightarrow{\epsilon(d_2)} s''$, for some $s' \in S$;

– (0-DELAY) for every $s, s' \in S$, $s \xrightarrow{\epsilon(0)} s'$ iff $s = s'$.

Following [Wan90], we now proceed to define versions of the transition relations that abstract away from the internal evolution of states as follows:

$$
\begin{array}{lll}
s \stackrel{a}{\Longrightarrow} s' & \text{iff} & \exists s''.\quad s \xrightarrow{\tau}{}^{*} s'' \xrightarrow{a} s' \\
s \stackrel{\epsilon(d)}{\Longrightarrow} s' & \text{iff} & \text{there exists a computation} \\
 & & s = s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \ldots \xrightarrow{\alpha_n} s_n = s' \quad (n \geq 0) \text{ where} \\
 & & (a) \quad \forall i \in \{1, .., n\}. \quad \alpha_i = \tau \text{ or } \alpha_i \in \mathcal{D} \\
 & & (b) \quad d = \sum \{d_i \mid \alpha_i = \epsilon(d_i)\}
\end{array}
$$

By convention, if the set $\{d_i \mid \alpha_i = \epsilon(d_i)\}$ is empty, then $\sum\{d_i \mid \alpha_i = \epsilon(d_i)\}$ is 0. With this convention, the relation $\stackrel{\epsilon(0)}{\Longrightarrow}$ coincides with $\xrightarrow{\tau}{}^{*}$, i.e., the reflexive, transitive closure of $\xrightarrow{\tau}$. Note that the derived transition relation $\stackrel{a}{\Longrightarrow}$ only abstracts from internal transitions *before* the actual execution of action a.

Definition 2.2. Let $\mathcal{T}_i = \langle \Sigma_i, \mathcal{L}, s_i^0, \longrightarrow_i \rangle$ $(i \in \{1, 2\})$ be two TLTSs. The parallel composition of $\mathcal{T}_1$ and $\mathcal{T}_2$ is the TLTS

$$\mathcal{T}_1 \parallel \mathcal{T}_2 = \langle \Sigma_1 \times \Sigma_2, \mathcal{D} \cup \{\tau\}, (s_1^0, s_2^0), \longrightarrow \rangle$$

where the transition relation $\longrightarrow$ is defined by the rules in Table 1. In Table 1, and in the remainder of the paper, we use the more suggestive notation $s \parallel s'$ in lieu of (s, s').

This definition of parallel composition forces the composed TLTSs to synchronize on actions (all but τ-actions) and delays, but with the particularity that delaying is only possible when no synchronization on actions is. This amounts to requiring that all actions in $\mathcal{A}$ be *urgent*. The reader familiar with TCCS [Wan90] may have noticed that the above definition of parallel composition precisely corresponds to a TCCS parallel composition in which all the actions in $\mathcal{A}$ are restricted upon. The use of this kind of parallel composition yields closed systems, of the type that can be analyzed using UPPAAL [BLL+95], and is inspired by the pioneering work by De Nicola and Hennessy on testing equivalence for processes [DNH84].

(1)	$\dfrac{s_1 \xrightarrow{\tau} s_1'}{s_1 \| s_2 \xrightarrow{\tau} s_1' \| s_2}$	(2)	$\dfrac{s_2 \xrightarrow{\tau} s_2'}{s_1 \| s_2 \xrightarrow{\tau} s_1 \| s_2'}$
(3)	$\dfrac{s_1 \xrightarrow{a} s_1' \qquad s_2 \xrightarrow{\overline{a}} s_2'}{s_1 \| s_2 \xrightarrow{\tau} s_1' \| s_2'}$		
(4)	$\dfrac{s_1 \xrightarrow{\epsilon(d)} s_1' \qquad s_2 \xrightarrow{\epsilon(d)} s_2'}{s_1 \| s_2 \xrightarrow{\epsilon(d)} s_1' \| s_2'}$		$\forall t \in [0, d[, a \in \mathcal{A}, s_1'', s_2''.$ $\neg(s_1 \xrightarrow{\epsilon(t)} s_1'' \xrightarrow{a} \ \wedge \ s_2 \xrightarrow{\epsilon(t)} s_2'' \xrightarrow{\overline{a}})$

where s_i, s_i', s_i'' are states of $\mathcal{T}_i$ $(i \in \{1,2\})$,
$a, \overline{a} \in \mathcal{A}$ and $d, t \in \mathbb{R}_{\geq 0}$.

Table 1: Rules defining the transition relation $\rightarrow$ in $\mathcal{T}_1 \parallel \mathcal{T}_2$

2.2 Timed Automata

Let C be a set of clocks. We use $\mathcal{B}(C)$ to denote the set of boolean expressions over atomic formulae of the form $x \sim p$, $x - y \sim p$, with $x, y \in C$, $p \in \mathbb{N}$, and $\sim \in \{<, >, =\}$. A *time assignment*, or *valuation*, v for C is a function from C to $\mathbb{R}_{\geq 0}$. For every time assignment v and $d \in \mathbb{R}_{\geq 0}$, we use $v + d$ to denote the time assignment which maps each clock $x \in C$ to the value $v(x) + d$. For every subset of clocks C', $[C' \rightarrow 0]v$ denotes the assignment for C which maps each clock in C' to the value 0 and agrees with v over $C \backslash C'$. Given a condition $g \in \mathcal{B}(C)$ and a time assignment v, the boolean value $g(v)$ describes whether g is satisfied by v or not.

Definition 2.3. A *timed automaton* is a tuple $A = \langle \mathcal{A}_\tau, N, n_0, C, E \rangle$ where N is a finite set of *nodes*, n_0 is the *initial node*, C is a finite set of *clocks*, and $E \subseteq N \times N \times \mathcal{A}_\tau \times 2^C \times \mathcal{B}(C)$ is a set of *edges*. The tuple $e = \langle n, n_e, \mu, r_e, g_e \rangle \in E$ stands for an edge from node n to node n_e (the *target* of e) with action μ, where r_e denotes the set of clocks to be reset to 0 and g_e is the enabling condition (or *guard*) over the clocks of A. For every node n and action μ, we use $E(n, \mu)$ to denote the set of edges emanating from n whose action is μ.

Example 2.4. The timed automaton depicted in Figure 1 has five nodes labelled n_0 to n_4, one clock x, and four edges. The edge from node n_1 to node n_2, for example, is guarded by $x \geq 0$, implies synchronization on a and resets clock x.

A *state* of a timed automaton A is a pair $\langle n, v\rangle$ where n is a node of A and v is a time assignment for C. The initial state of A is $\langle n_0, v_0\rangle$ where n_0 is the initial node of A and v_0 is the time assignment mapping all clocks in C to 0.

The operational semantics of a timed automaton A is given by the TLTS $\mathcal{T}_A = \langle \Sigma, \mathcal{L}, \sigma^0, \longrightarrow\rangle$, where Σ is the set of states of A, σ^0 is the initial state of A, and $\longrightarrow$ is the transition relation defined as follows:

$$\langle n, v\rangle \xrightarrow{\mu} \langle n', v'\rangle \text{ iff } \exists r, g.\langle n, n', \mu, r, g\rangle \in E \wedge g(v) \wedge v' = [r \to 0]v$$

$$\langle n, v\rangle \xrightarrow{\epsilon(d)} \langle n', v'\rangle \text{ iff } n = n' \text{ and } v' = v + d$$

where $\mu \in \mathcal{A}_\tau$ and $\epsilon(d) \in \mathcal{D}$.

Example 2.5. The following is a valid sequence of transitions for the timed automaton of Figure 1, where the number in brackets corresponds to the time assignment of clock x:

$$\langle n_0, \{0\}\rangle \xrightarrow{\tau} \langle n_1, \{0\}\rangle \xrightarrow{\epsilon(3.14)} \langle n_1, \{3.14\}\rangle \xrightarrow{a} \langle n_2, \{0\}\rangle \ .$$

3 Testing Automata

In this section we take the first steps towards the definition of *model checking via testing* by defining *testing*. Informally, testing involves the parallel composition of the tested automaton with a *test automaton*. The testing process then consists in performing reachability analysis in the composed system. We say that the tested automaton fails the test if a special reject state of the test automaton is reachable in the parallel composition from their initial configurations, and passes otherwise.

The formal definition of testing then involves the definition of what a test automaton is, how the parallel composition is performed and when the test has failed or succeeded. We now proceed to make these notions precise.

Definition 3.1. A *test automaton* is a tuple $T = \langle \mathcal{A}, N, N_T, n_0, C, E\rangle$ where $\mathcal{A}$, N, n_0, C, and E are as in Definition 2.3, and $N_T \subseteq N$ is the set of *reject nodes*.

Intuitively, a test automaton T interacts with a tested system, represented by a TLTS, by communicating with it. The dynamics of the interaction between the tester and the tested system is described by the parallel composition of the TLTS that is being tested and of $\mathcal{T}_T$. We now define failure and success of a test as follows.

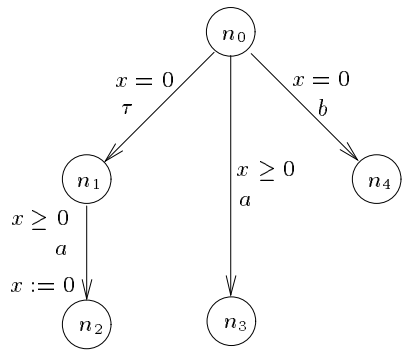

Figure 1: Timed automaton A

Definition 3.2. Let $\mathcal{T}$ be a TLTS and T be a test automaton.

- We say that a node n of T is reachable from a state $s_1 \parallel s_2$ of $\mathcal{T} \parallel \mathcal{T}_T$ iff there is a sequence of transitions leading from $s_1 \parallel s_2$ to a state whose $\mathcal{T}_T$ component is of the form $\langle n, u \rangle$.
- We say that a state s of $\mathcal{T}$ *fails the T-test* iff a reject node of T is reachable in $\mathcal{T} \parallel \mathcal{T}_T$ from the state $s \parallel \langle n_0, u_0 \rangle$, where $\langle n_0, u_0 \rangle$ is the initial state of $\mathcal{T}_T$. Otherwise, we say that s *passes the T-test.*

In the remainder of the paper, we shall mostly apply test automata to the TLTSs that give operational semantics to timed automata. In that case, we shall use the suggestive notation $A \parallel T$ in lieu of $\mathcal{T}_A \parallel \mathcal{T}_T$.

Example 3.3. Consider the timed automaton A of Figure 1 and the test automaton T_b of Figure 2(b). The reject node m_T of the test automaton is reachable from the initial state of $A \parallel T_b$, as follows:

1. first the automaton A can execute the τ-transition and go to node n_1, thus preempting the possibility of synchronizing on channel b with T,
2. now both automata can let time pass, thus enabling the τ-transition from node m_0 in T_b and making m_T reachable.

In this case we say that A fails the test. If we test A using the automaton T_a of Figure 2(a), then in all cases A and T_a must synchronize on a and no initial delay is possible. It follows that the reject node m_T of T_a is unreachable, and A passes the test.

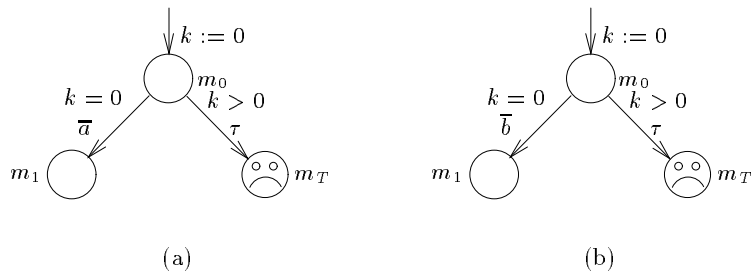

Figure 2: The test automata T_a and T_b

4 Safety Modal Logic

We consider a dense-time logic with clocks, which is a fragment of the one presented in [LLW95] and is suitable for the specification of safety and bounded liveness properties of TLTSs.

Definition 4.1. Let K be a set of clocks, disjoint from C. The set SBLL of (safety and bounded liveness) formulae over K is generated by the following grammar:

$$\begin{array}{lcl} \varphi & ::= & \mathtt{tt} \mid \mathtt{ff} \mid c \mid \varphi_1 \wedge \varphi_2 \mid c \vee \varphi \mid \mathbb{W}\varphi \mid \\ & & [a]\varphi \mid \langle a \rangle \mathtt{tt} \mid x \ \underline{\mathtt{in}}\ \varphi \mid X \mid \mathtt{max}(X, \varphi) \\ c & ::= & x \sim p \mid x - y \sim p \end{array}$$

where $a \in \mathcal{A}$, $x, y \in K$, $p \in \mathbb{N}$, $\sim \in \{<, >, =\}$, X is a formula variable and $\mathtt{max}(X, \varphi)$ stands for the maximal solution of the recursion equation $X = \varphi$.

Notation. For a set of formula clocks $\{y_1, \ldots, y_n\}$ and a formula φ, we write $\{y_1, \ldots, y_n\}\ \underline{\mathtt{in}}\ \varphi$ as a short-hand for $y_1\ \underline{\mathtt{in}}\ (y_2\ \underline{\mathtt{in}}\ \cdots (y_n\ \underline{\mathtt{in}}\ \varphi) \cdots)$. If $n = 0$, then, by convention, $\emptyset\ \underline{\mathtt{in}}\ \varphi$ stands for φ.

A *closed recursive formula* of SBLL is a formula in which every formula variable X appears within the scope of some $\mathtt{max}(X, \varphi)$ construct. In the remainder of this paper, every formula will be closed, unless specified otherwise.

Given a TLTS $\mathcal{T} = \langle S, \mathcal{L}, s^0, \longrightarrow \rangle$, we interpret the closed formulae in SBLL over extended states. An *extended state* is a pair $\langle s, u \rangle$ where s is a state of $\mathcal{T}$ and u is a time assignment for the formula clocks in K.

Definition 4.2. Consider a TLTS $\mathcal{T} = \langle S, \mathcal{L}, s^0, \longrightarrow \rangle$. The satisfaction relation $\models_w$ is the largest relation satisfying the implications in Table 2.

We say that $\mathcal{T}$ weakly satisfies φ, written $\mathcal{T} \models_w \varphi$, when $\langle s^0, u_0 \rangle \models_w \varphi$, where u_0 is the time assignment mapping every clock in K to 0. In the sequel, for a timed automaton A, we shall write $A \models_w \varphi$ in lieu of $\mathcal{T}_A \models_w \varphi$.

The weak satisfaction relation is closed with respect to the relation $\xrightarrow{\tau}^*$, in the sense of the following proposition.

Proposition 4.3. *Let $\mathcal{T} = \langle S, \mathcal{L}, s^0, \longrightarrow \rangle$ be a TLTS. Then, for every $s \in S$, $\varphi \in$ SBLL and valuation u for the clocks in K, $\langle s, u \rangle \models_w \varphi$ iff, for every s' such that $s \xrightarrow{\tau}^* s'$, $\langle s', u \rangle \models_w \varphi$.*

The reader familiar with the literature on variations on Hennessy-Milner logic [Mil89] and on its real-time extensions [Wan91] may have noticed that our definition of the satisfaction relation is rather different from the standard one presented in the literature. For instance, one might expect the clause of the definition of the satisfaction relation for the formula $\langle a \rangle \mathtt{tt}$ to read

$$\langle s, u \rangle \models_w \langle a \rangle \mathtt{tt} \quad \text{implies} \quad s \stackrel{a}{\Longrightarrow} s' \text{ for some } s' \ . \tag{1}$$

Recall, however, that our main aim in this paper is to develop a logical specification language for timed automata for which the model checking problem can be effectively reduced to deciding reachability. More precisely, for every formula $\varphi \in$ SBLL, we aim at constructing a test automaton T_φ such that every extended state $\langle s, u \rangle$ of a timed automaton satisfies φ iff it passes the test T_φ (in a sense to be made precise in Defn. 5.1). With this aim in mind, a reasonable proposal for a test automaton for the formula $\langle a \rangle \mathtt{tt}$, interpreted as in (1), is the automaton depicted in Figure 2(a). However, it is not hard to see that such an automaton could be brought into its reject node m_T by one of its possible interactions with the timed automaton associated with the TCCS agent $a + \tau$. This is due to the fact that, because of the definition of parallel composition we have chosen, a test automaton cannot prevent the tested state from performing its internal transition leading to a state where an a-action is no longer possible. (In fact, it is not too hard to generalize these ideas to show that *no* test automaton for the formula $\langle a \rangle \mathtt{tt}$ exists under the interpretation given in (1).) Similar arguments may be applied to all the formulae in the logic SBLL that involve occurrences of the modal operator $[a]$ and/or of the primitive proposition $\langle a \rangle \mathtt{tt}$.

$$
\begin{array}{rcl}
\langle s,u\rangle \models_w \mathtt{tt} & \Rightarrow & true \\
\langle s,u\rangle \models_w \mathtt{ff} & \Rightarrow & false \\
\langle s,u\rangle \models_w c & \Rightarrow & c(u) \\
\langle s,u\rangle \models_w \varphi_1 \wedge \varphi_2 & \Rightarrow & \forall s'.\ s \xrightarrow{\tau}{}^{*} s' \text{ implies } \langle s',u\rangle \models_w \varphi_1 \text{ and } \langle s',u\rangle \models_w \varphi_2 \\
\langle s,u\rangle \models_w c \vee \varphi & \Rightarrow & \forall s'.\ s \xrightarrow{\tau}{}^{*} s' \text{ implies } c(u) \text{ or } \langle s',u\rangle \models_w \varphi \\
\langle s,u\rangle \models_w [a]\varphi & \Rightarrow & \forall s'.\ s \stackrel{a}{\Longrightarrow} s' \text{ implies } \langle s',u\rangle \models_w \varphi \\
\langle s,u\rangle \models_w \langle a\rangle\mathtt{tt} & \Rightarrow & \forall s'.\ s \xrightarrow{\tau}{}^{*} s' \text{ implies } s' \xrightarrow{a} s'' \text{ for some } s'' \\
\langle s,u\rangle \models_w \mathbb{V}\varphi & \Rightarrow & \forall d \in \mathbb{R}_{\geq 0} \forall s'.\quad s \stackrel{\epsilon(d)}{\Longrightarrow} s' \text{ implies } \langle s',u+d\rangle \models_w \varphi \\
\langle s,u\rangle \models_w x\ \underline{\mathtt{in}}\ \varphi & \Rightarrow & \forall s'.\quad s \xrightarrow{\tau}{}^{*} s' \text{ implies } \langle s',[\{x\}\to 0]u\rangle \models_w \varphi \\
\langle s,u\rangle \models_w \mathtt{max}(X,\varphi) & \Rightarrow & \forall s'.\quad s \xrightarrow{\tau}{}^{*} s' \text{ implies } \langle s',u\rangle \models_w \varphi\{\mathtt{max}(X,\varphi)/X\}
\end{array}
$$

Table 2: Weak satisfaction implications

The reader might have also noticed that the language SBLL only allows for a restricted use of the logical connective 'or'. This is due to the fact that it is impossible to generate test automata even for simple formulae like $\langle a\rangle\mathtt{tt} \vee [b]\mathtt{ff}$—cf. Propn. 5.4.

Notation. Given a state $\langle n, v\rangle$ of a timed automaton, and a valuation u for the formula clocks in K, we write $\langle n, v : u\rangle$ for the resulting extended state.

Example 4.4. Assume that y is the only formula clock. For the automaton of Figure 1 the following holds:

- $\langle n_0, \{x\} \to 0 : \{y\} \to 0\rangle \not\models_w \langle b\rangle\mathtt{tt}$;
- $\langle n_0, \{x\} \to 0 : \{y\} \to 0\rangle \models_w \mathbb{W}\langle a\rangle\mathtt{tt}$.

5 Model checking via testing

In Sect. 3 we have seen how we can perform tests on timed automata. We now aim at using test automata to determine whether a given timed automaton weakly satisfies a formula in L. As already mentioned, this approach to model checking for timed automata is not merely a theoretical curiosity, but it is the way in which model checking of properties other than plain reachability ones is routinely carried out in a verification tool like UPPAAL. In order to achieve our goal, we shall define a "compilation" procedure to obtain a test automaton from the formula we want to test for. By means of this compilation procedure, we automate the process of

generating test automata from logical specifications—a task which has so far required a high degree of ingenuity and is error-prone.

Definition 5.1. Let φ be a formula in SBLL and T_φ be a test automaton over clocks $\{k\} \cup K$, k fresh.

- For every extended state $\langle n, v : u\rangle$ of a timed automaton A, we say that $\langle n, v : u\rangle$ passes the T_φ-test iff no reject node of T_φ is reachable from the state $\langle n, v\rangle \parallel \langle m_0, \{k\} \to 0 : u\rangle$, where m_0 is the initial node of T_φ.
- We say that the test automaton T_φ *weakly tests* for the formula φ iff the following holds: for every timed automaton A and every extended state $\langle n, v : u\rangle$ of A, $\langle n, v : u\rangle \models_w \varphi$ iff $\langle n, v : u\rangle$ passes the T_φ-test.

Theorem 5.2. *For every closed formula φ in* SBLL*, there exists a test automaton T_φ that weakly tests for it.*

Proof. (Sketch.) The test automata are constructed by structural induction on open formulae. (The Uppaal implementation of the constructions is depicted in Figures 3 and 4.) It can be shown that, for every closed formula φ, the resulting automaton T_φ weakly tests for φ. The details of the proof will be presented in the full version of the paper.

Corollary 5.3. *Let A be a timed automaton. Then, for every $\varphi \in$* SBLL*, there exists a test automaton T_φ with a reject node m_T such that $A \models_w \varphi$ iff node m_T is not reachable in $A \parallel T_\varphi$.*

As remarked in Sect. 4, the logic SBLL only allows for a restricted use of the 'or' operator. This is justified by the following negative result.

Proposition 5.4. *The formula $\langle a\rangle \mathtt{tt} \vee [b]\mathtt{ff}$ is not weakly testable.*

Proof. (Sketch.) Assume, towards a contradiction, that a test automaton T weakly tests for the formula $\langle a\rangle \mathtt{tt} \vee [b]\mathtt{ff}$. Then the timed automaton associated with the TCCS agent b must fail the T-test. Using the assumption that T weakly tests for $\langle a\rangle \mathtt{tt} \vee [b]\mathtt{ff}$, by a careful analysis of an arbitrary computation leading to the reject node of T in $b \parallel T$, we infer that such a computation must involve one b-synchronization preceded by zero or more τ-transitions from the tester T. It follows that a reject node in T can also be reached in $(a + b) \parallel T$. As $a + b$ weakly satisfies the formula $\langle a\rangle \mathtt{tt} \vee [b]\mathtt{ff}$, this contradicts the assumption that T weakly tests for $\langle a\rangle \mathtt{tt} \vee [b]\mathtt{ff}$.

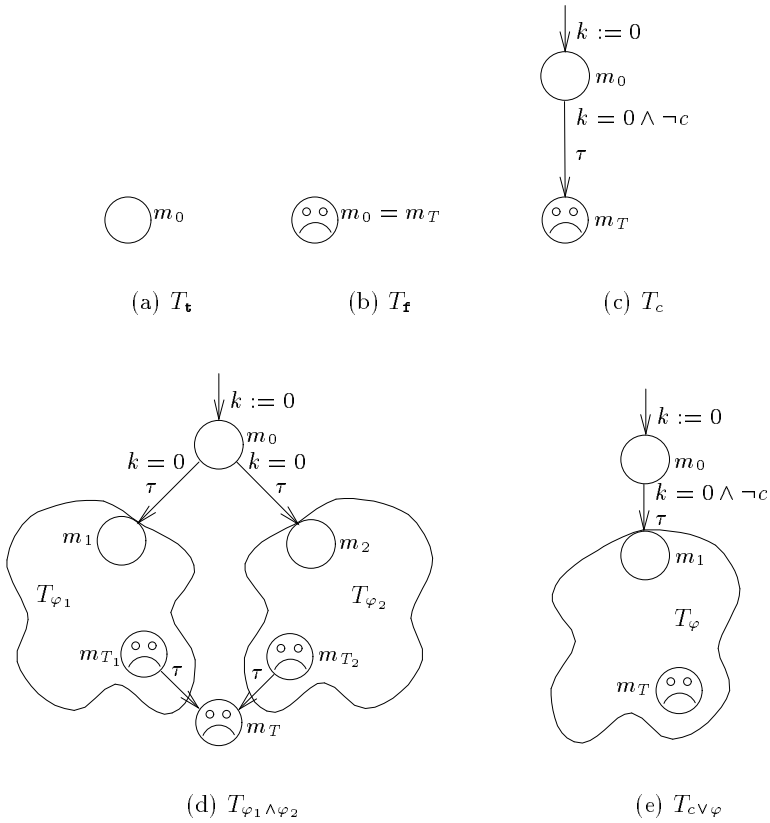

(a) $T_{\mathtt{t}}$ (b) $T_{\mathtt{f}}$ (c) T_c

(d) $T_{\varphi_1 \wedge \varphi_2}$ (e) $T_{c \vee \varphi}$

Figure 3: Test automata for SBLL sub-formulae

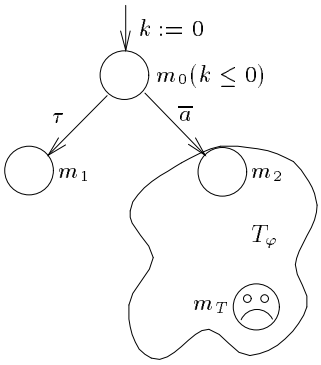

(a) $T_{[a]\varphi}$

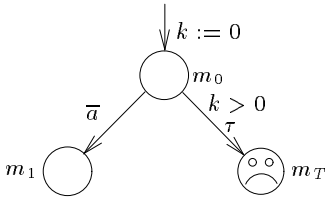

(b) $T_{\langle a \rangle \mathtt{t}}$

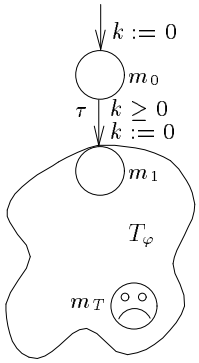

(c) $T_{\forall\varphi}$

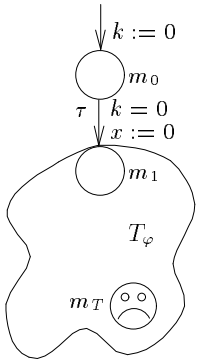

(d) $T_{x\ \underline{\mathtt{in}}\ \varphi}$

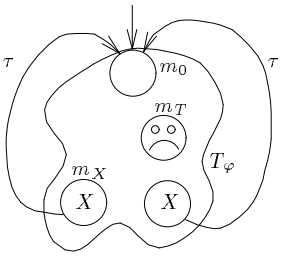

(e) $T_{\mathtt{max}(X,\varphi)}$

Figure 4: Test automata for SBLL sub-formulae (cont.)

6 Implementation in UPPAAL

The UPPAAL *constructs* The implementation of testing using the parallel composition operator presented in Sect. 3 requires a model of communicating timed automata with *urgent actions* (cf. rule (4) in Table 1). This feature is available in the UPPAAL model. The test automata are inductively obtained from the formula in a constructive manner, according to the constructions shown in Figures 3 and 4. In these constructions all actions in $\mathcal{A}$ are intended to be urgent. As in UPPAAL it is not possible to guard edges labelled with urgent actions, the theoretical construction for $T_{[a]\varphi}$ used in the proof of Thm. 5.2 is implemented by means of node invariants.

Simplification of the test automaton In certain cases, it is possible to optimize the construction of a test automaton from a formula by applying heuristics. Here we just remark on two possible simplifications. One is with respect to $T_{\varphi_1 \wedge \varphi_2}$ (Figure 3(d)) and the other one with respect to $T_{x\ \underline{\mathtt{in}}\ \varphi}$ (Figure 4(d)). Both simplifications involve the elimination of the τ-transitions emanating from node m_0. This leads to the constructs shown in Figures 5(a) and 5(b). The test automaton of Figure 5(a) is obtained by setting the initial nodes of T_{φ_1} and T_{φ_2} to be the same node m_0, and the same for the reject node m_T. For $T_{x\ \underline{\mathtt{in}}\ \varphi}$, the reset $x := 0$ is added to the incoming edge of T_φ. Nevertheless, these simplifications cannot be applied in the following cases:

$T_{\varphi_1 \wedge \varphi_2}$ When the *and* operator involves the conjunction of

- $[a]\varphi$ and $\langle a \rangle \mathtt{tt}$, or
- $[a]\varphi$ and $\mathbb{W}\varphi$, or
- $\langle a \rangle \mathtt{tt}$ and $\mathbb{W}\varphi$

the proposed simplification leads to incorrect test automata. This is because there is a different interpretation of evolving time in each operand, by, for example, leading to a reject state in one operand and to a safe one in the other one, or simply not being allowed in one case and being necessary in the other.

$T_{x\ \underline{\mathtt{in}}\ \varphi}$ The $\underline{\mathtt{in}}$ operator can be simplified only when it is not an operand in an *and* operation which has already been simplified. This is because the reset of the variable would affect the other branch of the *and* construct.

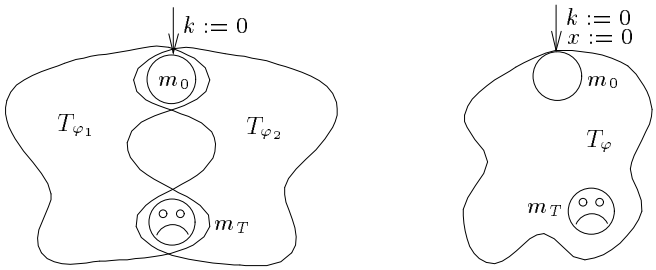

(a) $T_{\varphi_1 \wedge \varphi_2}$ (b) $T_{x\ \underline{\mathtt{in}}\ \varphi}$

Figure 5: New simplified constructs

High level operators The basic constructs of the logic SBLL can be used to define high level temporal operators, which may be used to simplify the writing of logical specifications (and substantiate our claim that SBLL can indeed express safety and bounded liveness properties). Here we confine ourselves to showing how to define the temporal operators `until`, `before` and `inv`:

$$\begin{aligned}
\varphi\ \mathtt{until}\ c &\stackrel{\text{def}}{=} \mathtt{max}(X, c \vee (\varphi \wedge \bigwedge_a [a]X \wedge \mathbb{W}X)) \\
\varphi\ \mathtt{until}_{\leq t}\ c &\stackrel{\text{def}}{=} x\ \ \underline{\mathtt{in}}\ ((\varphi \wedge x \leq t)\ \mathtt{until}\ c) \\
\mathtt{before}_t\ c &\stackrel{\text{def}}{=} \mathtt{tt}\ \mathtt{until}_{\leq t}\ c \\
\mathtt{inv}\ \varphi &\stackrel{\text{def}}{=} \mathtt{max}(X, \varphi \wedge \bigwedge_a [a]X \wedge \mathbb{W}X)\ \ .
\end{aligned}$$

The intuitive meanings of the above temporal operators are as follows: φ `until` c is true iff no matter how long the systems delays or what action transitions it takes, φ holds at least until c holds; φ $\mathtt{until}_{\leq t}$ c is its time bounded version, meaning that φ must hold at least until c holds, and c must hold within t time units; $\mathtt{before}_t$ c is true iff c is true within t time units; finally, `inv` φ is true iff no matter how long the systems delays or what action transitions it takes, φ always holds.

Optimized UPPAAL implementations of test automata for these constructs are shown in Figures 9 and 10, which may be found at the end of the paper. The above defined constructs express intuitively clear properties of real-time systems. However, as witnessed by the constructions in Figures 9 and 10, the definition of appropriate test automata for them is

a nontrivial task, which may be beyond the average user of a verification tool like UPPAAL. The compilation of formulae into test automata developed in this paper, and implemented in UPPAAL is a first step towards making model-checking technology more accessible to actual designers of real-time systems.

7 Example

Consider a number of stations connected on an Ethernet-like medium, following a basic CSMA/CD protocol as the one considered in [JLS96]. On top of this basic protocol, we want to design a protocol without collisions (applicable for example to real time plants). In particular, we want to guarantee an upper bound on the transmission delay of a buffer, assuming that the medium does not lose or corrupt data, and that the stations function properly. The simplest solution is to introduce a dedicated master station which asks the other stations whether they want to transmit data to another station (see Figure 6). Such a master station has to take into account the possible buffer delays within the receiving stations to ensure that the protocol enjoys the following properties: (1) collision cannot occur, (2) the transmitted data eventually reach their destination, (3) data which are received have been transmitted by a sender, and (4) there is a known upper bound on the transmission delay, assuming error-free transmission.

Modelling and verification of such a protocol in UPPAAL has been presented in [JLS96], where the details of such a modelling may be found. Here we only focus on the external view of the behaviour of the system. The observable actions are: user i sending a message, written `send_i!`, and user j receiving a message, written `recv_j!`, for $i, j = \{1, 2, 3\}$. The verification of the protocol presented in *op. cit.* was based on the ad hoc generation of test automata from informal specifications of system requirements. Indeed, some of the test automata that resulted from the informal requirements were rather complex, and it was difficult to extract their semantic interpretation. We now have at our disposal a precise property language to *formally* describe the expected behaviour of the protocol, together with an automatic compilation of such specifications into test automata, and we can therefore apply the aforementioned theory to test the behaviour of the protocol.

One of the requirements of the protocol is that there must be an upper bound on the transmission delay. Assuming that this upper bound is 4, this property can be expressed by means of the following formula in

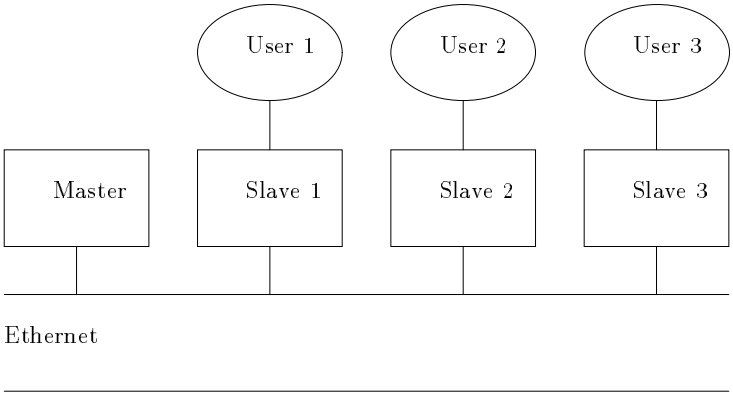

Figure 6: The Ethernet

SBLL:

$$\texttt{inv}\ \Big([\texttt{send_1!}]s\ \underline{\texttt{in}}\ \mathbb{W}\big([\texttt{recv_2!}](s < 4) \wedge [\texttt{recv_3!}](s < 4)\big)\Big)$$

This formula states that it invariantly holds that whenever user 1 sends a message, it will be received by users 2 and 3 within 4 units of time. Note that we consider transmission to be error-free, so the message will eventually be received. What we are interested in is the delay expressed by clock s. The test automaton corresponding to this formula is shown in Figure 7. (Note that, although the formula above expresses the required behaviour of the protocol in a very direct way, its encoding as a test automaton is already a rather complex object—which we were glad not to have to build by hand!)

In a similar way, the following properties represent the upper bounds between any two sending actions of user 1 and user 2, and between any two consecutive sending actions of user 1:

$$\texttt{inv}\ \Big([\texttt{send_1!}]s\ \underline{\texttt{in}}\ \mathbb{W}[\texttt{send_2!}](s < 5)\Big)$$
$$\texttt{inv}\ \Big([\texttt{send_1!}]s\ \underline{\texttt{in}}\ \mathbb{W}[\texttt{send_1!}](s < 17)\Big)$$

In order to experiment with our current implementation of the test automata construction in UPPAAL, we have also carried out the verification of several other properties of the protocol. For instance, we have verified that, under the assumption that the master waits for two time units before sending out its enquiries, the protocol has a round-trip time bound

of 18 time units, and that no faster round-trip exists. However, we have verified that changing the waiting time in the master to zero will allow for faster round-trip times. The details of these experiments will be reported in the full version of this study.

8 Characteristic Formulae and Ready Simulation

In the verification of realistic reactive systems, it is often useful to replace the individual components of the system under verification with more abstract versions before building the model of the complete system. This abstraction must, of course, be carried out in such a way that every property enjoyed by the resulting abstract model should also hold of the original, more detailed system description. (Cf. [TB97] for an impressive recent example of this general strategy applied to the verification of a high bandwidth communication chip.) In this section, we shall show how the results developed in this paper can be used to support this type of hierarchical approach to verification. More precisely, we shall show how the logic SBLL can be used to define characteristic properties [SI94] for nodes of τ-free, deterministic timed automata with respect to a timed version of the ready simulation preorder [LS91, BIM95]. As τ-free, deterministic timed automata are prime candidates for use as abstractions of more complex systems, the use of characteristic formulae allows us to formally, and automatically, justify abstractions using the model checking algorithm via reachability testing we have presented in the previous sections, and implemented in UPPAAL. The timed version of the ready simulation preorder that we shall consider is defined as follows:

Definition 8.1. Let $\mathcal{T} = \langle S, \mathcal{L}, s^0, \longrightarrow \rangle$ be a TLTS. We define the preorder $\preceq_w^{RS}$ as the largest binary relation over S such that if $s_1 \preceq_w^{RS} s_2$, then

1. whenever $s_1 \stackrel{a}{\Longrightarrow} s_1'$, then $s_2 \stackrel{a}{\Longrightarrow} s_2'$ for some s_2' such that $s_1' \preceq_w^{RS} s_2'$;
2. whenever $s_1 \stackrel{\epsilon(d)}{\Longrightarrow} s_1'$, then $s_2 \stackrel{\epsilon(d)}{\Longrightarrow} s_2'$ for some s_2' such that $s_1' \preceq_w^{RS} s_2'$; and
3. if $s_2 \stackrel{a}{\longrightarrow} s_2'$ for some s_2', then $s_1 \stackrel{a}{\longrightarrow} s_1'$ for some s_1'.

The main usage that we envisage for the relation $\preceq_w^{RS}$ is in justifying abstraction steps in verification. To this end, we expect that if $s_1 \preceq_w^{RS} s_2$ holds, then every property of the abstract state s_2 is also a property of s_1. This is the import of the following result.

Theorem 8.2. *Assume that $s_1 \preceq_w^{RS} s_2$. Then, for every $\varphi \in$ SBLL and valuation u for the formula clocks, $\langle s_2, u \rangle \models_w \varphi$ implies $\langle s_1, u \rangle \models_w \varphi$.*

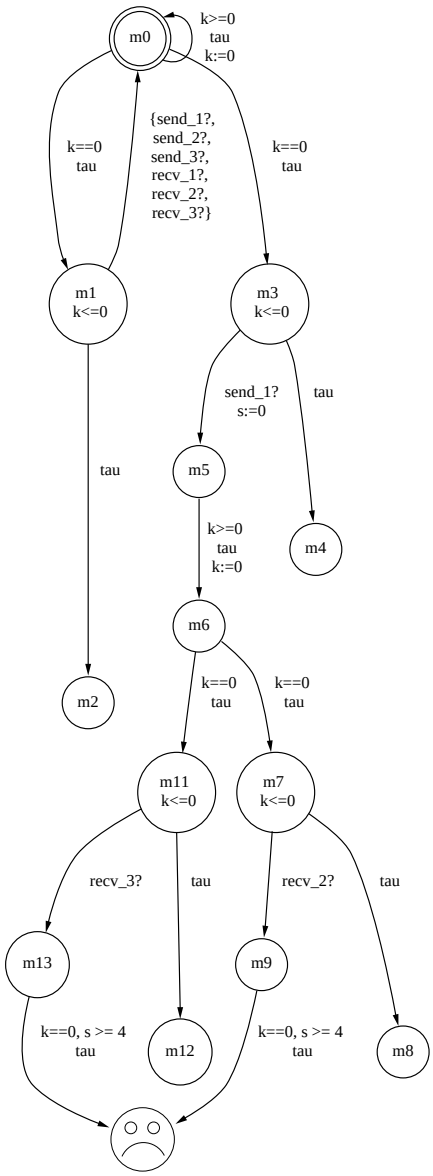

Figure 7: Test automaton for the invariant property

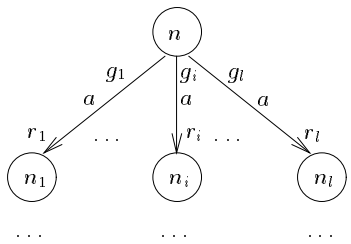

Figure 8: Node n of a timed automaton and its a-successors.

Consider the portion of a general test automaton shown in Figure 8. In the figure we can see the nodes that are reachable from node n by an a-transition (g_i represents the guard in the edge leading to node n_i and r_i the clocks to be reset in that edge). When it is the case that for every node n and for every action a, the guards g_i are disjoint, i.e. $g_i \wedge g_j$ is unsatisfiable when $i \neq j$, then the timed automaton A is *deterministic*.

We now proceed to define the characteristic formula for the nodes of a τ-free, deterministic timed automaton with respect to the timed ready simulation preorder introduced above. For the sake of clarity, in the following definition we shall specify recursive formulae using recursion equations in lieu of the $\mathtt{max}(X, \varphi)$ construct.

Definition 8.3. Let A be a τ-free, deterministic timed automaton. For every node n of A, we define the characteristic formula $\phi(n)$ as follows:

$$\phi(n) \stackrel{\text{def}}{=} \mathbb{W}\Big(\bigwedge_a [a]\Big(\textsc{Enabled}(n,a) \wedge \textsc{Match}(n,a)\Big) \wedge \bigwedge_a \textsc{Out}(n,a)\Big) \tag{2}$$

where

$$\begin{aligned}
\textsc{Enabled}(n,a) &\stackrel{\text{def}}{=} \bigvee_{e \in E(n,a)} g_e \\
\textsc{Match}(n,a) &\stackrel{\text{def}}{=} \bigwedge_{e \in E(n,a)} (g_e \Rightarrow r_e \ \underline{\mathtt{in}}\ \phi(n_e)) \\
\textsc{Out}(n,a) &\stackrel{\text{def}}{=} \textsc{Enabled}(n,a) \Rightarrow \langle a \rangle \mathtt{tt} \ .
\end{aligned}$$

Let A be a timed automaton with initial node n_0. We define the characteristic formula of A, notation $\phi(A)$, to be $\phi(n_0)$.

Intuitively, the formula $\phi(n)$ requires that, no matter how much a state s delays, and no matter how an action a is performed, then

- there should be at least one a-labelled edge of n that is enabled by the current value of the clocks (formula $\textsc{Enabled}(n, a)$);
- the successor state of s satisfies the characteristic formula of the target of the only a-labelled edge of n that is enabled, modulo the appropriate resets of clocks (formula $\textsc{Match}(n, a)$ and determinism of n); and
- s has an a-labelled transition if an a-labelled edge is enabled in n by the current value of the clocks (formula $\textsc{Out}(n, a)$).

These intuitive remarks capture the essence of the proof of the following result.

Theorem 8.4. *Let A and B be two timed automata. Assume that B is deterministic and τ-free. Then, for every node n of A and m of B, and valuations v, w,*

$$\langle n, v\rangle \preceq_w^{RS} \langle m, w\rangle \quad \textit{iff} \quad \langle n, v : w\rangle \models_w \phi(m) \ .$$

Corollary 8.5. *Given two timed automata A and B, A with initial state $\langle n_0, v_0\rangle$ and B deterministic, without τ transitions and with initial state $\langle m_0, w_0\rangle$, then $A \preceq_w^{RS} B$ iff $A \models_w \phi(m_0)$.*

Remark. The characteristic formula for the timed simulation preorder can be obtained by simply omitting the sub-formula involving $\textsc{Out}(n, a)$ from (2). Both the characteristic formula constructions can be extended to timed automata with node invariants.

As a corollary of these results, we obtain that timed ready simulation is "testable" in the sense of this paper. In particular, we have shown how the problem of checking the existence of a behavioural relation between states of two timed automata can be recast as a reachability problem that can be efficiently handled by Uppaal. We envisage that such an approach can, for instance, be applied to yield automatic tool support for the justification of the abstraction steps used in, e.g., [TB97]. In order to take full advantage of this approach, abstraction steps need to be justified using a precongruence relation with respect to the chosen notion of parallel composition. Here we just remark that neither timed simulation nor timed ready simulation is preserved by TCCS parallel composition—which is the one adopted in Uppaal to combine open systems. However, both the aforementioned relations are preserved by TCCS parallel composition if the more abstract system is τ-free. These are precisely the abstraction steps supported by our method.

9 Concluding Remarks

As argued in, e.g., [Wol], efficient algorithms for deciding reachability questions can be used to tackle many common problems related to verification. In this study, following the lead of [VW86], we have shown how to reduce model-checking of safety and bounded liveness properties expressible in the real-time logic SBLL to checking for reachability of reject states in suitably constructed test automata. This approach allows us to take full advantage of the core of the computational engine of the tool UPPAAL [BLL+95], which consists of a collection of efficient algorithms that can be used to perform reachability analysis over timed automata. As the logic that we consider is powerful enough to describe characteristic properties [SI94] for nodes of timed automata with respect to (ready) simulation, our approach to model-checking also allows us to reduce the computation of behavioural relations to reachability analysis. Historically, model-checking and reachability analysis were amongst the first problems shown to be decidable for timed automata [ACD90]. The decidability of behavioural equivalences and preorders was shown at a later date in [Cer92]. This study may be seen as tracing back the decidability of a behavioural relation, viz. (ready) simulation, to that of the reachability problem via model-checking.

The practical applicability of the approach to model-checking that we have developed in this paper has been tested on a basic CSMA/CD protocol. More experimental activity will be needed to fully test the feasibility of model-checking via reachability testing. So far, all the case studies carried out with the use of UPPAAL (see, e.g., [BGK+96, JLS96, KP95]) seem to support the conclusion that this approach to model-checking can indeed be applied to realistic case studies, but further evidence needs to be accumulated to substantiate this claim. In this process of experimentation, we also expect to further develop a collection of heuristics that can be used to reduce the size of the test automata obtained by means of our automatic translation of formulae into automata.

In this study, we have shown how to translate the formulae in the logic SBLL into test automata in such a way that model-checking can be reduced to testing for reachability of distinguished reject nodes in the generated automata. Indeed the logic presented in this study is remarkably close to being completely expressive with respect to reachability properties. In fact, as it will be shown in a companion paper [ABBL97], a slight extension of the logic considered here allows us to reduce any reachability property for a composite system $S \parallel T$ to a model-checking problem of

S.

The interpretation of the formulae in our specification logic presented in Table 2 abstracts from the internal evolution of real-time processes in a novel way. A natural question to ask is whether the formulae in the logic SBLL are testable, in the sense of this paper, when interpreted with respect to the transition relation $\longrightarrow$. In the full version of this work, we shall show that this is indeed possible if the test automata are allowed to have *committed locations* [BGK+96], and the definition of the parallel composition operator is modified to take the nature of these locations into account. We expect, however, that the weak interpretation of the logic will be more useful in practical applications of our approach to model-checking.

The results that we have developed show that a timed version of ready simulation is testable, in the sense of this paper. This conclusion seems to be in agreement with the analysis of behavioural relations carried out in [AV93] within the framework of quantales. Whether our results can be justified by means of a general theory à la Abramsky and Vickers is an interesting topic for further theoretical research. It would also be interesting to investigate the connections between our investigations and the seminal study [BIM95], where ready simulation is characterized as the largest precongruence, with respect to all the GSOS definable operations, which is contained in the preorder induced by completed trace inclusion.

Acknowledgements. We thank Patricia Bouyer for her help in the implementation of the tool, and Kåre Jelling Kristoffersen for his proof-reading.

References

ABBL97. L. Aceto, P. Bouyer, A. Burgueño, and K. G. Larsen. The limitations of testing for timed automata, 1997. Forthcoming paper.

ACD90. R. Alur, C. Courcoubetis, and Dill D. Model-checking for real-time systems. In *Proc. of the 5th. Annual Symposium on Logic in Computer Science, LICS'90*, pages 41–425. IEEE Computer Society Press, 1990.

AD94. R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:183–235, 1994.

AH94. R. Alur and T.A. Henzinger. A really temporal logic. *Journal of the ACM*, 41(1):181–204, 1994. Preliminary version appears in Proc. 30th FOCS, 1989.

AV93. S. Abramsky and S. Vickers. Quantales, observational logic and process semantics. *Mathematical Structures in Computer Science*, 3(2):161–227, June 1993.

BGK+96. J. Bengtsson, D. Griffioen, K. Kristoffersen, K. G. Larsen, F. Larsson, P. Pettersson, and W. Yi. Verification of an audio protocol with bus collision using UPPAAL. In R. Alur and T. A. Henzinger, editors, *Proc. of the 8th. International Conference on Computer-Aided Verification, CAV'96*, volume 1102 of *Lecture Notes in Computer Science*, New Brunswick, New Jersey, USA, July 31 – August 3 1996. Springer-Verlag.

BIM95. B. Bloom, S. Istrail, and A. R. Meyer. Bisimulation can't be traced. *Journal of the Association for Computing Machinery*, 42(1):232–268, January 1995.

BLL+95. J. Bengtsson, K. G. Larsen, F. Larsson, P. Pettersson, and W. Yi. UPPAAL - a tool suite for automatic verification of real-time systems. In *Proc. of the 4th DIMACS Workshop on Verification and Control of Hybrid Systems*, New Brunswick, New Jersey, 22–24 October 1995.

Cer92. K. Cerans. Decidability of bisimulation equivalence for parallel timer processes. In *Proc. 4th Workshop on Computer-Aided Verification*, Montreal, Canada, 1992.

CS91. R. Cleaveland and B. Steffen. Computing behavioural relations, logically. In J. L. Albert, B. Monien, and M. Rodriguez Artalejo, editors, *Proc. of the 18th International Colloquium on Automata, Languages and Programming, ICALP'91*, volume 510 of *Lecture Notes in Computer Science*, pages 127–138, Madrid, Spain, July 8–12 1991. Springer-Verlag.

DNH84. R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 34:83–133, 1984.

HHWT95. T. A. Henzinger, P.-H. Ho, and H. Wong-Toi. HyTech: the next generation. In *Proc. of the 16th Real-time Systems Symposium, RTSS'95*. IEEE Computer Society press, 1995.

HK94. T. A. Henzinger and P.W. Kopke. Undecidability results for hybrid systems. In *Proc. of the Workshop on Hybrid Systems and Autonomous Control*, Ithaca, NY, October 1994. Also appeared as Cornell University technical report TR95-1483.

HKPV95. T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya. What's decidable about hybrid automata? In *Proc. of the 27th Annual ACM Symposium on Theory of Computing, STOC'95*, pages 373–382, 1995. Also appeared as Cornell University technical report TR95-1541.

HM85. M. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the Association for Computer Machinery*, 32(1):137–161, January 1985.

HNSY94. T. A. Henzinger, X. Nicollin, J. Sifakis, and S. Yovine. Symbolic model checking for real-time systems. *Information and Computation*, 111(2):193–244, 1994.

HWT95. P.-H. Ho and H. Wong-Toi. Automated analysis of an audio control protocol. In P. Wolper, editor, *Proc. of the 7th. International Conference on Computer-Aided Verification, CAV'95*, volume 939 of *Lecture Notes in Computer Science*, pages 381–394, Lige, Belgium, July 1995. Springer-Verlag.

JLS96. H. E. Jensen, K. G. Larsen, and A. Skou. Modelling and analysis of a collision avoidance protocol using SPIN and UPPAAL. In *DIMACS Workshop SPIN '96, 2nd International SPIN Verification Workshop on Algorithms, Applications, Tool Use, Theory*. Rutgers University, New Jersey, USA, 1996.

KP95. K.J. Kristoffersen and P. Pettersson. Modelling and analysis of a steam generator using UPPAAL. In *Proc. of the 7th Nordic Workshop on Programming Theory*, Göteborg, Sweden, November 1–3 1995.

LLW95. F. Laroussinie, K. G. Larsen, and C. Weise. From timed automata to logic - and back. In J. Wiedermann and P. Hájek, editors, *Proc. of the 20th. International Symposium on Mathematical Foundations of Computer Science, MFCS'95*, volume 969 of *Lecture Notes in Computer Science*, pages 529–539, Prague, Czech Republic, August 28 - September 1 1995. Springer-Verlag.

LS91. K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, September 1991.

Mil89. R. Milner. *Communication and Concurrency*. Series in Computer Science. Prentice Hall International, 1989.

OY93. A. Olivero and S. Yovine. *Kronos: a tool for verifying real-time systems. User's guide and reference manual*. VERIMAG, Grenoble, France, 1993.

Pnu77. A. Pnueli. The temporal logic of programs. In *Proc. of the 19th Annual Symposium on Foundations of Computer Science, FOCS'77*, pages 46–57. IEEE Computer Society, 1977.

SI94. B. Steffen and A. Ingólfsdóttir. Characteristic formulae for processes with divergence. *Information and Computation*, 110(1):149–163, April 1994.

SVD97. J. Springintveld, F. W. Vaandrager, and P. D'Argenio. Testing timed automata. Technical Report CSI-R9712, Katholieke Universiteit Nijmegen, Computing Science Institute, Faculty of Mathematics and Informatics, August 1997.

TB97. S. Taşiran and R. K. Brayton. STARI: A case study in compositional verification and hierarchical timing verification. In Orna Grumberg, editor, *Proc. of the 9th. International Conference on Computer-Aided Verification, CAV'97*, volume 1254 of *Lecture Notes in Computer Science*, pages 191–201, Haifa, Israel, June 22–25 1997. Springer-Verlag.

Var96. M. Y. Vardi. An automata-theoretic approach to linear temporal logic. In F. Moller and G. Birtwistle, editors, *Logics for Concurrency: structure versus Automata*, volume 1043 of *Lecture Notes in Computer Science*, Berlin, 1996. Springer-Verlag.

VW86. M. Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *Proc. of the 1st. Annual Symposium on Logic in Computer Science, LICS'86*, pages 322–331. IEEE Computer Society Press, 1986.

VW94. M. Y. Vardi and P. Wolper. Reasoning about infinte computations. *Information and Computation*, 115:1–37, 1994.

Wan90. Y. Wang. Real-time behaviour of asynchronous agents. In J.C.M. Baeten and J.W. Klop, editors, *Proc. of the Conference on Theories of Concurrency: Unification and Extension, CONCUR'90*, volume 458 of *Lecture Notes in Computer Science*, pages 502–520, Amsterdam, The Netherlands, August 27–30 1990. Springer-Verlag.

Wan91. Y. Wang. *A calculus of real time systems*. PhD thesis, Chalmers university of Technology, Göteborg, Sweden, 1991.

Wol. P. Wolper. Where could SPIN go next? A unifying approach to exploring infinite state spaces. Slides for an invited talk at the *1997 SPIN Workshop*, Enschede, The Netherlands. Available at the URL `http://www.montefiore.ulg.ac.be/~pw/papers/psfiles/SPIN4-97.ps`.

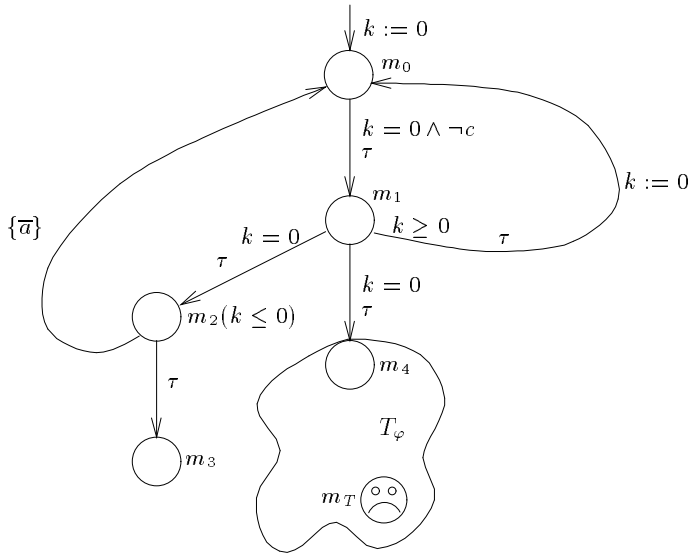

(a) $T_{\varphi \texttt{ until } c}$

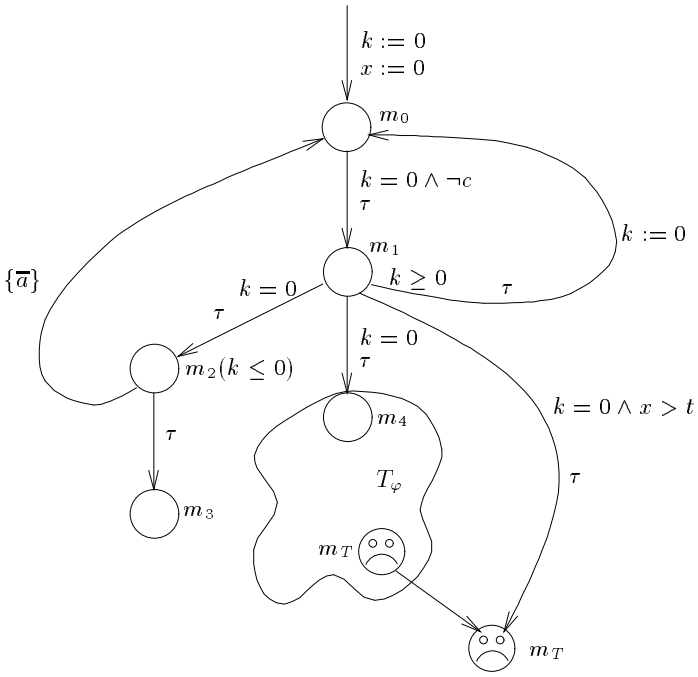

(b) $T_{\varphi \texttt{ until}_{\leq t} c}$

Figure 9: High level operators

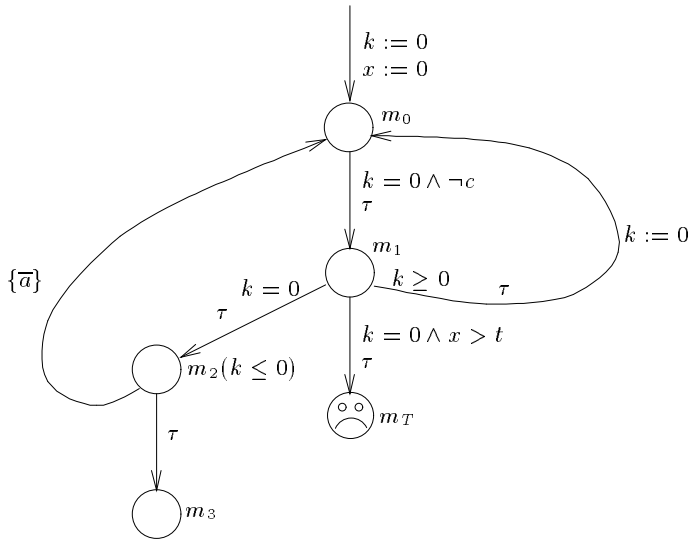

(a) $T_{\texttt{before}_t\ c}$

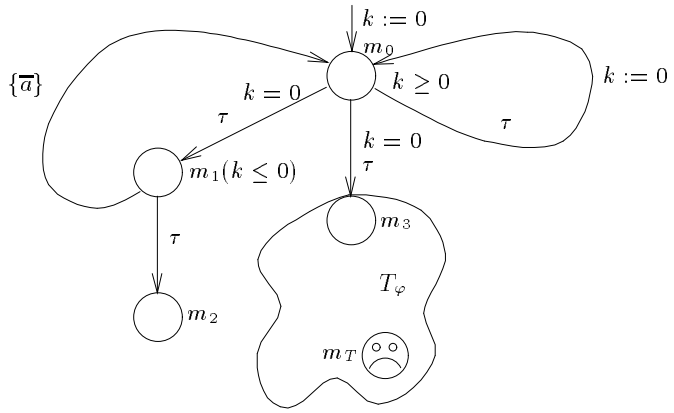

(b) $T_{\texttt{inv}\ \varphi}$

Figure 10: High level operators (cont.)

Recent BRICS Report Series Publications

RS-97-29 **Luca Aceto, Augusto Burgueño, and Kim G. Larsen.** ***Model Checking via Reachability Testing for Timed Automata*****. November 1997. 29 pp.**

RS-97-28 **Ronald Cramer, Ivan B. Damgård, and Ueli Maurer.** ***Span Programs and General Secure Multi-Party Computation*****. November 1997. 27 pp.**

RS-97-27 **Ronald Cramer and Ivan B. Damgård.** ***Zero-Knowledge Proofs for Finite Field Arithmetic or: Can Zero-Knowledge be for Free?*** **November 1997. 33 pp.**

RS-97-26 **Luca Aceto and Anna Ingólfsdóttir.** ***A Characterization of Finitary Bisimulation*****. October 1997. 9 pp. To appear in** ***Information Processing Letters*****.**

RS-97-25 **David A. Mix Barrington, Chi-Jen Lu, Peter Bro Miltersen, and Sven Skyum.** ***Searching Constant Width Mazes Captures the AC^0 Hierarchy*****. September 1997. 20 pp. To appear in** ***STACS '98: 15th Annual Symposium on Theoretical Aspects of Computer Science Proceedings*****, LNCS, 1998.**

RS-97-24 **Søren B. Lassen.** ***Relational Reasoning about Contexts*****. September 1997. 45 pp. To appear as a chapter in the book** ***Higher Order Operational Techniques in Semantics*****, eds. Andrew D. Gordon and Andrew M. Pitts, Cambridge University Press.**

RS-97-23 **Ulrich Kohlenbach.** ***On the Arithmetical Content of Restricted Forms of Comprehension, Choice and General Uniform Boundedness*****. August 1997. 35 pp.**

RS-97-22 **Carsten Butz.** ***Syntax and Semantics of the logic $\mathcal{L}^{\lambda}_{\omega\omega}$*****. July 1997. 14 pp.**

RS-97-21 **Steve Awodey and Carsten Butz.** ***Topological Completeness for Higher-Order Logic*****. July 1997. 19 pp.**

RS-97-20 **Carsten Butz and Peter T. Johnstone.** ***Classifying Toposes for First Order Theories*****. July 1997. 34 pp.**