

A test generation framework for *quiescent* real-time systems

Laura Brandán Briones and Ed Brinksma

Faculty of Computer Science, University of Twente,
P.O.Box 217, 7500AE Enschede,
The Netherlands. Fax - (31 53)-489-3247
{brandanl,brinksma}@cs.utwente.nl

Abstract. We present an extension of Tretmans' theory and algorithm for test generation for input-output transition systems to real-time systems. Our treatment is based on an operational interpretation of the notion of *quiescence* in the context of real-time behaviour. This gives rise to a family of implementation relations parameterized by observation durations for *quiescence*. We define a nondeterministic (parameterized) test generation algorithm that generates test cases that are sound with respect to the corresponding implementation relation. The test generation is also exhaustive in the sense that for each non-conforming implementation a test case can be generated that detects the non-conformance.

1 Introduction

Although testing has always been the most important technique for the validation of software systems it has only become a topic of serious academic research in the past decade or so. In this period research on the use of formal methods for model-driven test generation and execution of functional test cases has led to a number of promising methods and tools for systematic black-box testing of systems, e.g. [AJR⁺99,JE03,JCCTC96,JCCTC97]. Most of these approaches are limited to the qualitative behaviour of systems, and exclude quantitative aspects such as real-time properties. The explosive growth of embedded software, however, has also caused a growing need to extend existing testing theories to the testing of real-time reactive systems. In this paper we present an extension of Tretmans' **ioco** theory for test generation [J.T96] for input-output transition systems that includes real-time behaviour.

A central concept in the non-timed theory is the notion of *quiescence*, which characterizes systems states that will not produce any output response without the provision of a new input stimulus. By treating *quiescence* as a special sort of system output the notion of behavioural trace can be generalized to include observations of *quiescence*. In turn, this leads to an implementation relation that defines unambiguously if implemented behaviour conforms to a given specification model, viz. if after all specified generalized traces of the implementation all possible generalized outputs are allowed according to the specification. Or, more informally, if all outputs and *quiescence* are correctly predicted by the specification.

In practice, the above implementation criterion means that implementations can be more deterministic than their specifications. Although it is good engineering practice not to introduce unnecessary nondeterminism in reactive systems, it is often unavoidable in the context of testing, and should therefore be part of a sensible testing theory. The reason for this is twofold:

- although the implementation under test may be deterministic, it can often only be tested through a testing environment that includes operating system features, communication media, etc. that typically introduce nondeterminism into the observed behaviour.

- an implementation under test often consists of concurrent components in an asynchronous parallel composition. The loss of information about the relative progress of components results in nondeterministic properties of their integrated behaviour.

Our proposed extension of the **ioco** theory to real-time systems is based on an operational interpretation of the notion of *quiescence*. This gives rise to a family of implementation relations parameterized by observation durations for *quiescence*. We define a nondeterministic (parameterized) test generation algorithm that generates test cases that are sound with respect to the corresponding implementation relation. This means that when an implementation fails any of the generated tests, must be non-conforming. The algorithm is also exhaustive in the sense that for every non-conforming implementation a test case can be generated that will detect its non-conformance.

The rest of this paper is organized as follows. Section 2 is a summary of ioco-testing. Section 3 introduces the model of timed input-output transition systems and our conformance relation. Section 4 presents the real-time test generation algorithm. Section 5 illustrates the theory with an example in the setting of timed automata. Section 6 compares our achievements to related work. Finally, section 7 presents the conclusions and future work.

2 Untimed ioco-testing

In this section we give a brief summary of the theory of input-output transition systems and the implementation relation **ioco**, from [J.T96]. An input-output transition system (IOTS) is a labelled transition system whose action set is partitioned into input actions, whose occurrence is controlled by the environment of the system, and output actions, whose occurrence is controlled by the system. We follow the convention that input actions are identified by names followed by a $?$ -symbol, and output actions by names followed by a $!$ -symbol.

Definition 1. *An Input-Output Transition System (IOTS) is a 4-tuple $\langle S, s_0, Act, \rightarrow \rangle$, where*

- S is a countable, non-empty set of states
- $s_0 \in S$ is the initial state
- Act is a countable set of labels, partitioned into input (Act_I) and output (Act_U) actions, with $Act_I \cap Act_U = \emptyset$
- $\rightarrow \subseteq (S \times (Act \cup \{\tau\}) \times S)$ is the transition relation.

To obtain the desired distribution of the control over the occurrence of input and output actions, an *IOTS* must be able to accept all input actions in any of its states, i.e. an *IOTS* is *input-enabled*. Symmetrically, the environment of an *IOTS* system is assumed always to be able to accept any system output action. Given that input actions are always enabled, traditional deadlock states cannot exist in *IOTS* systems. In this setting a weaker notion becomes relevant: states that cannot produce output actions, and that can only be left by (further) supply of input actions: such states are called *quiescent states*.

Definition 2. *Let p be a state in an IOTS over Act , then*

$$p \text{ is quiescent, denoted by } \delta(p) \text{ iff } \forall \mu \in (Act_U \cup \{\tau\}) : p \not\stackrel{\mu}{\rightarrow}$$

where $p \not\stackrel{\mu}{\rightarrow}$ denotes that there is no transition from p labelled with μ .

The idea is to treat *quiescence* as an observable event. This can be formalized by extending the transition relation as follows: a transition $p \xrightarrow{\delta} p$ is added for each *quiescent* state p such that:

$$\delta(p) \text{ iff } p \xrightarrow{\delta} p.$$

We can now introduce the notion of a suspension trace, i.e. a sequence of input actions, output actions and *quiescences* as they can occur when observing an *IOTS*.

Definition 3. Let the transition relation (\rightarrow) be extended with quiescence transitions $p \xrightarrow{\delta} p$. Then, the suspension traces of process p are:

$$\text{Straces}(p) \stackrel{\text{def}}{=} \{\sigma \in \text{Act}_\delta^* \mid p \xRightarrow{\sigma}\}$$

where Act_δ means $\text{Act} \cup \{\delta\}$.

We use the well-known notation $p \xRightarrow{\sigma}$ to say that there exists a state q that is reachable from p by performing σ abstracting from the internal actions (see next section for a formal definition of the \Rightarrow -relation).

For input-output transition systems all output actions that are enabled in state p , including the *quiescence* action δ , are collected into the set $\text{out}(p)$.

Definition 4. Let p be a state and let P be a set of states in a input-output transition system, then

$$\begin{aligned} \text{out}(p) &\stackrel{\text{def}}{=} \{\mu \in \text{Act}_U \mid p \xrightarrow{\mu}\} \cup \{\delta \mid \delta(p)\} \\ \text{out}(P) &\stackrel{\text{def}}{=} \bigcup_{p \in P} \text{out}(p) \end{aligned}$$

The correctness of an *input-enabled* implementation $i \in \text{IOTS}$ with respect to a specification $s \in \text{IOTS}$ is now given as an *implementation relation*, the **io**co relation, defined by:

$$i \text{ io}co s \stackrel{\text{def}}{=} \forall \sigma \in \text{Straces}(s) : \text{out}(i \text{ after } \sigma) \subseteq \text{out}(s \text{ after } \sigma)$$

where $(p \text{ after } \sigma)$ denotes the set of states that is reachable from p by performing the *Strace* σ .

Informally, this means that the implementation i is **io**co-correct with respect to a specification s , if and only if, after all possible behaviours of the specification ($\forall \sigma \in \text{Straces}(s)$), any output action μ produced by the implementation ($\mu \in \text{out}(i \text{ after } \sigma)$) can also occur as an output of the specification ($\mu \in \text{out}(s \text{ after } \sigma)$). In particular, this should also hold for the special action *quiescence* (δ), which models the absence of outputs.

Using this theory, the following algorithm for test derivation is presented in [J.T96]. Here, a test case is understood as the specification of the behaviour of a deterministic and finite testing process that can be carried out with an implementation under test. Also the behaviour of test cases can be described by input-output transition systems, where the occurrence of a δ -action¹ in a test case corresponds to the detection of *quiescence* in an implementation, i.e.


¹ In [J.T96] the action symbol θ is used for the observation of *quiescence*. We prefer to use δ for both *quiescence* and its observation, in line with the philosophy that identical actions synchronize.

the observation that no output is produced. In practice, the observation of δ is implemented using a time-out of sufficiently long duration.

A state of a test case is either a terminal state labelled (**fail** or **pass**), or a state that offers one particular input (to the implementation), or a state accepting all possible outputs (from the implementation), including the δ -action. The class of test cases is denoted as $TEST$.

Let $s \in IOTS$ be a specification with initial state s_0 and S be a non-empty set of states, with initially $S = \{s_0\}$ (S represents the set of all possible states in which the implementation can be at the current stage of the test case execution). As in [J.T96], we present the test using a syntax inspired by LOTOS [ISO89]; the concrete semantics for the timed case will be explained in section 4.

The algorithm for the generation of test cases $t \in TEST$ from S consists of a finite number of recursive applications of a nondeterministic selection between one of the following three alternatives:



1. The single state test case **pass**, which stops the recursion in the algorithm and thus terminates the test case:

$$t := \mathbf{pass}$$

2. Test case t supplies the input μ and behaves as test case t' :

$$t := \mu; t'$$

where $\mu \in Act_I$, $(S \mathbf{after} \mu) \neq \emptyset$, and t' is obtained by recursively applying the algorithm to $(S \mathbf{after} \mu)$.

3. Test case t checks the next output of the implementation; if it is a valid response the test case continues recursively; if it is an invalid response, i.e. $\mu \notin out(S)$ then the test case terminates in **fail**. The observation of *quiescence* δ is treated separately:

$$\begin{aligned} t := & \Sigma\{\mu_i; t_i \mid \mu_i \in Act_U \wedge \mu_i \in out(S)\} \\ & + \Sigma\{\delta; t_\delta \mid \delta \in out(S)\} \\ & + \Sigma\{\mu_j; \mathbf{fail} \mid \mu_j \in Act_U \wedge \mu_j \notin out(S)\} \\ & + \Sigma\{\delta; \mathbf{fail} \mid \delta \notin out(S)\} \end{aligned}$$

where t_i and t_δ are obtained by recursively applying the algorithm for $(S \mathbf{after} \mu_i)$ and $(S \mathbf{after} \delta)$, respectively.

This algorithm is implemented in the test generation tool TORX. Indeed, the developed theory plus its tool support provided by TORX have proven quite a useful and a successful approach to the functional testing of reactive systems [JE03].

Example 5. This is an adaptation of an example due to Rom Langerak [R.L90]. Figure 1 shows two coffee machines with peculiar behaviour. They illustrate the importance of being able to observe *quiescence* and act on the basis of this information, as formalized by the **ioco**

relation. Suppose both graphs are saturated with input action transitions in each state by adding input self-loops for all input transitions that are not explicitly given. After accepting a $coin?$ both machines can end up in state q_1 (q_2) where only the the action $tea?$ ($coffee?$) leads to the desired effect, viz. the corresponding output of $tea!$ ($coffee!$). In the machine on the left the user can switch between the tea - and $coffee$ -mode by $bang?$ -ing the machine after noticing that $coffee?$ ($tea?$) does not produce $coffee!$ ($tea!$). In the machine on the right such switches are not possible. These two machines are distinguished by the **io** relation, because:

$$out(\text{left-hand machine after } coin? \cdot coffee? \cdot \delta \cdot bang? \cdot coffee?) = \{coffee!\}$$

whereas

$$out(\text{right-hand machine after } coin? \cdot coffee? \cdot \delta \cdot bang? \cdot coffee?) = \{\delta\}$$

However, if the special action δ were be dropped from the definition of suspension traces and *out*-sets, leading to the input/output testing pre-order [J.T96], both machines would be indistinguishable because we could not use the *quiescence* of states q_1 and q_2 .

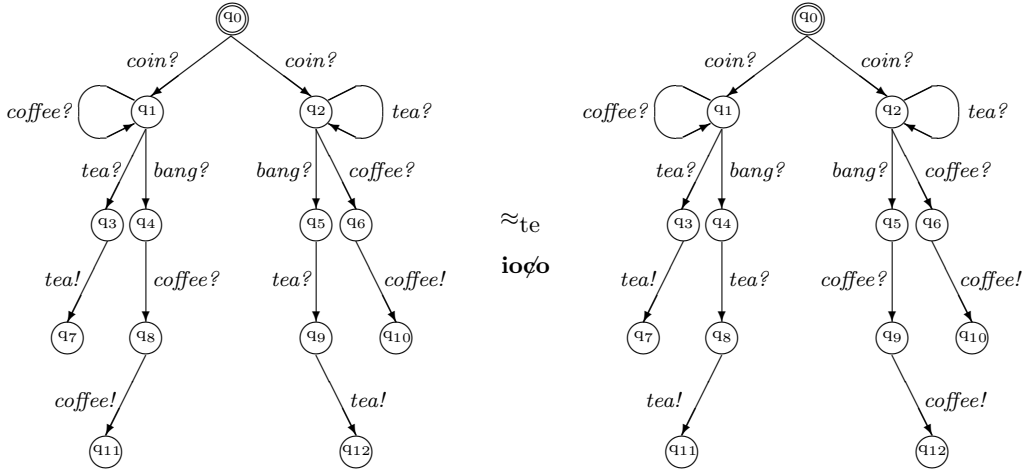


Fig. 1. The quirky coffee machine, a modification of [R.L90]

3 Implementation relations for real-time *quiescence*

3.1 Timed input-output transition systems

In this section we introduce the concept of Timed labelled transition systems, their properties and notation, and then specialize them to obtain the model of Timed input-output transition systems. After that, we proceed to obtain a conformance relation between a specification and an implementation, defined as timed input-output transition systems, analogous to the **io** relation for the untimed case.

We distinguish three types of actions: *time-passage actions*, *visible labelled actions* and the special *internal action* τ . All except the time-passage actions are thought of as occurring instantaneously, i.e. without consuming time. To specify time, a dense time domain is used, viz. the nonnegative reals (\mathbb{R}^+); no lower *a priori* bounds are imposed on the delays between events.

Definition 6. A *Timed labelled transition system (TLTS)* is a 4-tuple $\langle S, s_0, Act_{\tau\varepsilon}, \rightarrow \rangle$, where

- S is a non-empty set of states
- $s_0 \in S$ is the initial state
- $Act_{\tau\varepsilon} \stackrel{def}{=} Act \cup \{\tau\} \cup \mathcal{D}$ are the actions Act including the internal action τ and time-passage actions; where \mathcal{D} is $\{\varepsilon(d) \mid d \in \mathbb{R}^+\}$
- $\rightarrow \subseteq (S \times Act_{\tau\varepsilon} \times S)$ is the transition relation with the following consistency constraints:
 - **Time Determinism** whenever $s \xrightarrow{\varepsilon(d)} s'$ and $s \xrightarrow{\varepsilon(d)} s''$ then $s' = s''$
 - **Time Additivity** $\forall s, s'' \in S \wedge \forall d_1, d_2 \geq 0 : (\exists s' \in S : s \xrightarrow{\varepsilon(d_1)} s' \xrightarrow{\varepsilon(d_2)} s'') \text{ iff } s \xrightarrow{\varepsilon(d_1+d_2)} s''$
 - **Null Delay** $\forall s, s' \in S : s \xrightarrow{\varepsilon(0)} s' \text{ iff } s = s'$.

The labels in Act_ε ($Act_\varepsilon \stackrel{def}{=} Act \cup \mathcal{D}$) represent the observable actions of a system, i.e. labelled actions and passage of time; the special label τ represents an unobservable internal action. A transition $(s, \mu, s') \in \rightarrow$ is denoted as $s \xrightarrow{\mu} s'$. A *computation* is a finite or infinite sequence of transitions:

$$s_0 \xrightarrow{\mu_1} s_1 \xrightarrow{\mu_2} s_2 \xrightarrow{\mu_3} \dots \xrightarrow{\mu_{n-1}} s_{n-1} \xrightarrow{\mu_n} s_n (\rightarrow \dots)$$

A *timed trace* captures the observable aspects of a computation; it is the sequence of observable actions. The set of all *finite* sequences of actions over Act_ε is denoted by Act_ε^* , while ϵ denotes the empty sequence. If $\sigma_1, \sigma_2 \in Act_\varepsilon^*$ then $\sigma_1 \cdot \sigma_2$ is the concatenation of σ_1 and σ_2 .

We denote the class of all timed labelled transition systems over Act by $TLTS(Act)$. Some additional notations and properties are introduced in the next definitions.

Definition 7. Let $p = \langle S, s_0, Act_{\tau\varepsilon}, \rightarrow \rangle$ be a $TLTS(Act)$ with $s, s', s_i \in S; d, d', e \in \mathbb{R}^+; \mu_i \in Act_{\tau\varepsilon}; \beta_i \in Act; \alpha_i \in Act_\varepsilon; \alpha \in Act_\varepsilon^*$, then

$$\begin{array}{ll}
 s \xrightarrow{\mu_1 \dots \mu_n} s' & \stackrel{def}{=} \exists s_0, \dots, s_n : s = s_0 \xrightarrow{\mu_1} s_1 \xrightarrow{\mu_2} \dots \xrightarrow{\mu_n} s_n = s' \\
 s \xrightarrow{\mu_1 \dots \mu_n} & \stackrel{def}{=} \exists s' : s \xrightarrow{\mu_1 \dots \mu_n} s' \\
 s \not\xrightarrow{\mu_1 \dots \mu_n} & \stackrel{def}{=} \nexists s' : s \xrightarrow{\mu_1 \dots \mu_n} s' \\
 s \xrightarrow{\varepsilon} s' & \stackrel{def}{=} s = s' \text{ or } s \xrightarrow{\tau \dots \tau} s' \\
 s \xrightarrow{\beta_i} s' & \stackrel{def}{=} \exists s_1, s_2 : s \xrightarrow{\varepsilon} s_1 \xrightarrow{\beta_i} s_2 \xrightarrow{\varepsilon} s' \\
 s \xrightarrow{\varepsilon(d)} s' & \stackrel{def}{=} (\exists s_1, s_2 : s \xrightarrow{\varepsilon} s_1 \xrightarrow{\varepsilon(d)} s_2 \xrightarrow{\varepsilon} s') \text{ or } (\exists s_1, d', d'' : d' + d'' = d : s \xrightarrow{\varepsilon(d')} s_1 \xrightarrow{\varepsilon(d'')} s') \\
 s \xrightarrow{\alpha_1 \dots \alpha_n} s' & \stackrel{def}{=} \exists s_0 \dots s_n : s = s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} s_n = s' \\
 s \xrightarrow{\alpha} & \stackrel{def}{=} \exists s' : s \xrightarrow{\alpha} s' \\
 s \not\xrightarrow{\alpha} & \stackrel{def}{=} \nexists s' : s \xrightarrow{\alpha} s'.
 \end{array}$$

We do not always distinguish between a timed labelled transition system and its initial state: if $p = \langle S, s_0, Act_{\tau\varepsilon}, \rightarrow \rangle$, we will often identify the process p with its initial state s_0 , e.g. we write $p \xrightarrow{\alpha}$ instead of $s_0 \xrightarrow{\alpha}$.

Definition 8.

- $ttraces(p) \stackrel{def}{=} \{\sigma \in Act_\varepsilon^* \mid p \xrightarrow{\sigma}\}$
- $init(p) \stackrel{def}{=} \{\mu \in Act_{\tau\varepsilon} \mid p \xrightarrow{\mu}\}$
- $der(p) \stackrel{def}{=} \{p' \mid \exists \sigma \in Act_\varepsilon^* : p \xrightarrow{\sigma} p'\}$
- $p \text{ after } \sigma \stackrel{def}{=} \{p' \mid p \xrightarrow{\sigma} p'\}$
- $P \text{ after } \sigma \stackrel{def}{=} \bigcup_{p \in P} (p \text{ after } \sigma)$, where P is a set of states
- p is deterministic if for all $\sigma \in Act_\varepsilon^* : (p \text{ after } \sigma)$ has at most one element. If $\sigma \in ttraces(p)$, then $(p \text{ after } \sigma)$ is overloaded to denote this element.

As an example, the above definitions applied to the left-hand side machine of Figure 1 give us: $init(q_9) = \{tea?, coffee?, bang?, tea!\}$ (remember that all states are saturated with all inputs), $der(q_2) = \{q_2, q_5, q_6, q_9, q_{10}, q_{12}\}$ and $(q_0 \text{ after } coin?) = \{q_1, q_2\}$.

In the context of timed systems there are some further important properties.

Definition 9. Let $p = \langle S, s_0, Act_{\tau\varepsilon}, \rightarrow \rangle$ be a $TLTS(Act)$, then p is **time divergent**: if for all $s \in S$ there exists an infinite computation σ from s with infinite cumulative delay:

$$\forall s \in S : \exists \sigma \in Act_{\tau\varepsilon}^\omega : \sigma = \mu_1 \cdot \mu_2 \cdot \mu_3 \cdots : s \xrightarrow{\sigma} \wedge \Sigma\{d_i \mid \mu_i = \varepsilon(d_i)\} = \infty$$

p has **Zeno behaviour**: if there exists a state $s \in S$ and an infinite computation from s with infinitely many non-delay actions and finite cumulative delay:

$$\exists s \in S : \exists \sigma \in Act_{\tau\varepsilon}^\omega : \sigma = \mu_1 \cdot \mu_2 \cdot \mu_3 \cdots : s \xrightarrow{\sigma} \wedge |\{i \mid \mu_i \neq \varepsilon(d_i)\}| = \infty \\ \wedge \Sigma\{d_i \mid \mu_i = \varepsilon(d_i)\} < \infty.$$

We assume that for all $p \in TLTS$ we are working with, p is time divergent, and does not have Zeno behaviour.

We now introduce Timed input-output transition systems ($TIO TS$) to model timed systems for which the set of actions can be partitioned into *output actions* and *input actions*. To do this properly, we formalize the notion of *input enabling*: if an input action is initiated by the environment, the system is always prepared to participate in such an interaction: all the inputs can always be accepted without letting time pass. Also, we want to exclude the possibility that the flow of time in a system can be blocked because the environment does not provide certain input actions, i.e. there must be no *forced inputs*.

Definition 10. A Timed input-output transition system ($TIO TS$) is a timed labelled transition system $\langle S, s_0, Act_{\tau\varepsilon}, \rightarrow \rangle$ with Act partitioned into input actions, Act_I , and output actions, Act_U , ($Act_I \cup Act_U = Act, Act_I \cap Act_U = \emptyset$), that has the properties of

weak input enabling:

$$\forall s \in S : \forall \mu \in Act_I : s \xrightarrow{\mu}$$

no forced inputs: iff for all $s \in S$ there exists an infinite computation σ from s containing no input actions and with infinite cumulative delay:

$$\forall s \in S : \exists \sigma \in (Act_U \cup \{\tau\} \cup \mathcal{D})^\omega : \sigma = \mu_1 \cdot \mu_2 \cdots : s \xrightarrow{\sigma} \wedge \Sigma\{d_i \mid \mu_i = \varepsilon(d_i)\} = \infty$$

The class of timed input-output transition systems with input actions in Act_I and output actions in Act_U is denoted by $TIO TS(Act_I, Act_U) \subseteq TLTS(Act_I \cup Act_U)$.

A timed trace σ is a sequence of actions and delays, e.g. $\sigma = a? \cdot \varepsilon(d_1) \cdot \varepsilon(d_2) \cdot b!$. Obviously, it would be more natural to avoid consecutive delays, as in $\sigma = a? \cdot \varepsilon(d_1 + d_2) \cdot b!$. Such traces could alternatively be written as sequences of actions with relative time stamps, viz. $\sigma = a?(0) \cdot b!(d_1 + d_2)$. This idea motivates the definition of *normalized timed traces*.

Definition 11. Let $\sigma \in Act_\varepsilon^*$, then

- σ is a normalized timed trace iff $\sigma \in (\mathcal{D} \cdot Act)^*$
- $ntraces(p) = \{\sigma \in (\mathcal{D} \cdot Act)^* \mid p \xrightarrow{\sigma}\}$
- for normalized timed traces $\sigma = \varepsilon(d_0) \cdot a_0 \cdot \varepsilon(d_1) \cdot a_1 \cdots \varepsilon(d_n) \cdot a_n$ we also write $\widehat{\sigma} = a_0(d_0) \cdot a_1(d_1) \cdots a_n(d_n)$.

If a timed trace begins with an action it can always be converted to a *normalized timed trace* by combining delays, or adding zero delays $\varepsilon(0)$ in the appropriate places. But if a timed trace ends with a delay, such as $\sigma = \varepsilon(d_0) \cdot a? \cdot \varepsilon(d_1) \cdot b! \cdot \varepsilon(d_2)$ then is not possible to interpret it as a *normalized timed trace*. The next lemma shows, however, that in the presence of input enabledness *normalized timed traces* preserve the information of timed traces.

Lemma 12. Let $p_1, p_2 \in TIOTS(Act_I, Act_U)$, then

$$ttraces(p_1) \subseteq ttraces(p_2) \text{ iff } ntraces(p_1) \subseteq ntraces(p_2)$$

Proof.

[\Rightarrow] Obvious, as $ntraces(p_1) \subseteq ttraces(p_1) \subseteq ttraces(p_2)$,

every normalized trace of p_1 is a trace of p_2 and thus a normalized trace of p_2

[\Leftarrow] Let $\sigma \in ttraces(p_1)$ then by input enabledness: $\sigma \cdot \mu \in ttraces(p_1)$ for any $\mu \in Act_I$,

let $\widehat{\sigma \cdot \mu}$ be the normalized version of $\sigma \cdot \mu$, i.e. $\widehat{\sigma \cdot \mu} \in ntraces(p_1)$ and

therefore $\widehat{\sigma \cdot \mu} \in ntraces(p_2)$

But then because of the density of the delay transition relation also $\sigma \cdot \mu \in ttraces(p_2)$

and therefore $\sigma \in ttraces(p_2)$. □

From now on we will not distinguish between a timed trace σ and its normalization $\widehat{\sigma}$ if it exist.

Similarly to Tretmans' work, we proceed to introduce the notion of *quiescence* in the timed setting. In the presence of time we define a *quiescent* state as one where the system is unable to produce an output immediately or in the future without receiving further input stimuli.

Definition 13. Let $p \in TIOTS(Act_I, Act_U)$. A state s of p is quiescent, denoted by $\delta(s)$, iff

$$\forall \mu \in Act_U : \forall d \in \mathbb{R}^+ : s \not\xrightarrow{\mu(d)}$$

As before in the untimed case, we can start out by representing *quiescence* as a special action δ ($\delta \notin Act \cup \{\tau\}$), and extending the timed transition relation of a *TIOTS* p to include self-loop transitions $s \xrightarrow{\delta} s$ iff s is a *quiescent* state. Moreover, let $\Delta(p)$ denotes the extended transition system of p that is obtained in this way.

3.2 Timed implementation relation

The extension of the timed transition relation allows us to define the following relation over *TIOTS*.

Definition 14. Let p and $q \in \text{TIOTS}(\text{Act}_I, \text{Act}_U)$, then

$$q \sqsubseteq_{\text{tiorf}} p \text{ iff } \text{nttraces}(\Delta(q)) \subseteq \text{nttraces}(\Delta(p)).$$

For specifications $s \in \text{TIOTS}$ the *quiescent* states can, in principle, be identified by analyzing the timed transition system, i.e. we can assume that $\Delta(s)$ is at our disposal. For implementations i , however, we only can detect *quiescence* by waiting for outputs. But we cannot wait forever, and therefore need to choose a maximal duration M . This motivates the following parameterized version of $\sqsubseteq_{\text{tiorf}}$ where σ can only appear after M time-units.

Definition 15. Let p and $q \in \text{TIOTS}$, then

$$q \sqsubseteq_{\text{tiorf}}^M p \text{ iff } \Delta_M(q) \subseteq \Delta_M(p)$$

where $\Delta_M(r) \stackrel{\text{def}}{=} \text{nttraces}(\Delta(r)) \cap (\mathcal{D} \cdot \text{Act} \cup \varepsilon(M) \cdot \delta)^*$.

The above definition takes only into account observations of *quiescence* that are made after a minimal delay of M time units. Naturally this definition implies a pre-order.

Lemma 16. If $M_1 < M_2$, then

$$\text{if } q \sqsubseteq_{\text{tiorf}}^{M_1} p \text{ then } q \sqsubseteq_{\text{tiorf}}^{M_2} p.$$

Proof. Assume $q \sqsubseteq_{\text{tiorf}}^{M_1} p$ and let σ be a trace in $\Delta_{M_2}(q)$.

Transform σ into σ' by replacing every subsequence of σ of the form $\delta(M_2)\mu(d)$ by $\delta(M_1)\mu(d + (M_2 - M_1))$, and possibly $\delta(M_2)$ by $\delta(M_1)$ if it is the last action of σ . σ' must also be a trace of $\Delta(q)$, and therefore of $\Delta_{M_1}(q)$. As $q \sqsubseteq_{\text{tiorf}}^{M_1} p$, it follows that σ' is a trace in $\Delta_{M_1}(p)$. By postponing all observations of δ by $M_2 - M_1$ we get $\sigma \in \Delta_{M_2}(p)$, i.e. $\Delta_{M_2}(q) \subseteq \Delta_{M_2}(p)$ or $q \sqsubseteq_{\text{tiorf}}^{M_2} p$. \square

This is not without consequences: in contrast to the untimed case, time delays can change the system state, which has interesting consequences, as shown by the modification of the quirky coffee machine example of Figure 1.

Example 17. Figure 2 shows two quirky coffee machines with time. Again, suppose both graphs are saturated with input action transitions in each state by adding input self-loops for all input transitions that are not explicitly given. Note that here after introducing *money?* we can switch between the coffee and tea modes. If we order *coffee?* and *bang?* fast enough we always will have coffee in the right-hand machine and some times in the left-hand machine, but if we *bang?* after waiting for the *quiescence* we will not notice the difference between machines. It follows from the one that cannot switch modes. This is a consequence of the fact that observing *quiescence* takes time. For simplicity, in the figure, we use $m?$ for money, $b?$ for bang, $c?, c!$ for coffee, and $t?, t!$ for tea. We also suppose that each action resets the clock x and that $k < M$ (for simplicity we used the representation of timed automata).




Fig. 2. The quirky coffee machine with time

The output set of a given state of a system in $TIOTS(Act_I, Act_U)$ consists of the time stamped output actions that are allowed from that state (abstracting from τ -actions), including δ -actions after a delay of M time-units.

Definition 18. Let p be a state of an (extended) timed transition system in $TIOTS(Act_I, Act_U)$, then:

$$out_M(p) = \{\mu(d) \mid \mu \in Act_U \wedge p \xRightarrow{\mu(d)}\} \cup \{\delta(M) \mid p \xRightarrow{\delta(M)}\}$$

and for P a set of states, then:

$$out_M(P) = \bigcup_{p \in S} out_M(p).$$

Lemma 19. Let p and $q \in TIOTS(Act_I, Act_U)$, then

$$q \sqsubseteq_{\text{tiorf}}^M p \text{ iff } \forall \sigma \in (\mathcal{D} \cdot Act \cup \varepsilon(M) \cdot \delta)^* : out_M(q \text{ after } \sigma) \subseteq out_M(p \text{ after } \sigma)$$

Proof.

[\Rightarrow] Let $\sigma \in (\mathcal{D} \cdot Act \cup \varepsilon(M) \cdot \delta)^*$, then

if $\sigma \notin \text{ntraces}(\Delta(q))$, then $out_M(q \text{ after } \sigma) = \emptyset$

if $\sigma \in \text{ntraces}(\Delta(q))$, then

$\forall \mu \in out_M(q \text{ after } \sigma) : \sigma \cdot \mu \in \Delta_M(q)$ and $\sigma \cdot \mu \in \Delta_M(p)$, then
 $\mu \in out_M(p \text{ after } \sigma)$

[\Leftarrow] Let $\sigma \in \Delta_M(q)$, i.e., $\sigma \in \text{ntraces}(\Delta(q)) \cap (\mathcal{D} \cdot Act \cup \varepsilon(M) \cdot \delta)^*$

then, $\Delta(q) \text{ after } \sigma \neq \emptyset$

using the *no forced input* property: $out_M(q \text{ after } \sigma) \neq \emptyset$

then, $out_M(p \text{ after } \sigma) \neq \emptyset$

which tell us that: $\sigma \in \text{ntraces}(\Delta(p)) \wedge \sigma \in (\mathcal{D} \cdot Act \cup \varepsilon(M) \cdot \delta)^*$

then, $\sigma \in \Delta_M(p)$

□

Finally, we are in position to define the relation we use to test real time systems: \mathbf{tioco}_M . For p and $q \in \mathbf{TIOOTS}(Act_I, Act_U)$, q will be \mathbf{tioco}_M to p if the set of outputs of q *after every normalized timed trace* σ of p including observations $\delta(M)$, is a subset of the outputs of p after the same σ .

Definition 20. Let p and $q \in \mathbf{TIOOTS}(Act_I, Act_U)$, then

$$q \sqsubseteq_{\mathbf{tioco}}^M p \text{ iff } \forall \sigma \in \Delta_M(p) : out_M(q \text{ after } \sigma) \subseteq out_M(p \text{ after } \sigma)$$

we also write $\sqsubseteq_{\mathbf{tioco}}^M$ as \mathbf{tioco}_M .

3.3 An operational model

To obtain an effective theory of *quiescence* in a timed setting we need more than stipulating that observing *quiescence* takes time. Since with physical implementations we can only observe absence of outputs over finite time intervals we must stipulate when such observations will be interpreted as *quiescence*.

Definition 21. Let q be a \mathbf{TIOOTS} and $M \in \mathbb{R}^+$, then

- a state s of q is M -quiescent iff $\forall s' \in (s \text{ after } \varepsilon(M)) : s'$ is quiescent
- q is M -quiescent iff all states s of q are M -quiescent.

In line with the above development we now want to formalize how *normalized timed traces* of \mathbf{TIOOTS} s may be enriched directly with δ -actions. Whenever the *normalized timed trace* allows an action with a delay of more than M time-units this creates a possibility to observe *quiescence*. For example, if $M = 4$ and $\sigma = a?(2) \cdot b?(5) \cdot c!(3)$ is an observed trace then it is also possible to observe $\sigma' = a?(2) \cdot \delta(4) \cdot b?(1) \cdot c!(3)$. We formalize the addition of δ -observations to *normalized timed traces* as a formal relation δ_M between (extended) *normalized timed traces*.

Definition 22. Let σ, σ' be normal form of $\sigma, \sigma' \in (\mathcal{D} \cdot (Act \cup \delta))^*$, then

- $\sigma \delta_M \sigma'$ iff

$$\exists \sigma_1, \sigma_2 : \exists \mu : \exists d \geq M : \sigma = \sigma_1 \cdot \mu(d) \cdot \sigma_2 \wedge \sigma' = \sigma_1 \cdot \delta(M) \cdot \mu(d - M) \cdot \sigma_2$$

- let Σ be a set of *normalized timed traces*, then

$$\delta_M(\Sigma) = \text{pref} \left(\bigcup_{\sigma \in \Sigma} \{ \sigma' \mid \sigma \delta_M^* \sigma' \} \right)$$

where $\text{pref}(S)$ is interpreted as the prefix-closure of a set of traces S and δ_M^* is the reflexive transitive closure of the relation δ_M .

If δ -actions are introduced in *normalized timed traces* on the basis observations of delays of (at least) M time units, we must check for consistency, i.e. we must have the property expressed in the following lemma.

Lemma 23. Let $q \in \mathbf{TIOOTS}$ be M -quiescent, then

$$\nexists \sigma \in \delta_M(\text{nttraces}(q)) : \exists \mu \in Act_U : \sigma = \sigma' \cdot \delta(M) \cdot \mu(d)$$

Proof. Suppose that such a $\sigma = \sigma' \cdot \delta(M) \cdot \mu(d)$ with $\mu \in Act_U$ does exist. Then $\sigma' \cdot \mu(d + M)$ is a trace of the system too. It follows there are reachable states $s' \in (q \text{ after } \sigma')$ and $s'' \in (s' \text{ after } \varepsilon(M))$ with $s'' \xrightarrow{\mu(d)}$, i.e. s'' is not quiescent. Contradiction. \square

Corollary 24. *Let $q \in \text{TIOTS}$ be M -quiescent, then*

$$\delta_M(\text{nttraces}(q)) = \Delta_M(q).$$

Proof.

[\Rightarrow] This proof is carried out by induction of the number of $\delta(M)$ occurrences in the traces of q . It is obvious that for δ -free traces $\sigma \in \delta_M(\text{nttraces}(q))$ iff $\sigma \in \text{nttraces}(q)$ iff $\sigma \in \Delta_M(q)$. Now suppose that σ is of the form $\sigma_1 \cdot \delta(M) \cdot \mu(d) \cdot \sigma_2$, with n occurrences of $\delta(M)$ in σ_1 and σ_2 δ -free. We will prove that such σ are in $\delta_M(\text{nttraces}(q))$ iff they are $\Delta_M(q)$.

$$\begin{aligned} & \sigma_1 \cdot \delta(M) \cdot \mu(d) \cdot \sigma_2 \in \delta_M(\text{nttraces}(q)) \\ \text{(by def. of } \delta_M) & \Leftrightarrow \sigma_1 \cdot \mu(d + M) \cdot \sigma_2 \in \delta_M(\text{nttraces}(q)) \\ \text{(ind.hyp. on } \sigma_1) & \Leftrightarrow \sigma_1 \cdot \mu(d + M) \cdot \sigma_2 \in \Delta_M(q) \\ \text{(skip } \delta\text{-observation)} & \Leftrightarrow \sigma_1 \cdot \delta(M) \cdot \mu(d) \cdot \sigma_2 \in \Delta_M(q) \end{aligned}$$

Clearly, we have that $\sigma \in \Delta_M(q)$ implies $\sigma \in \delta_M(\text{nttraces}(q))$.

[\Leftarrow] Let $\sigma \in \delta_M(\text{nttraces}(q))$, then by the above implications that $\sigma = \sigma_1 \cdot \mu(d + M) \cdot \sigma_2$ is in $\Delta_M(q)$. Now suppose that $\sigma_1 \cdot \delta(M) \cdot \mu(d) \cdot \sigma_2 \notin \Delta_M(q)$. This means that there is a non-quiescent state in $(q \text{ after } \sigma_1 \cdot \varepsilon(M))$, i.e. there is an output action $\mu \in \text{Act}_U$ with $\sigma_1 \cdot \mu(d' + M) \in \Delta_M(q)$. Again, by the above implication we have $\sigma_1 \cdot \delta(M) \cdot \mu(d') \in \delta_M(\text{nttraces}(q))$. Because of the preceding lemma, this contradicts the M -quiescence of q .

If σ is not of the required format, it must be of the form $\sigma_1 \delta(M)$, and by input-enabledness we can extend it to the form $\sigma_1 \cdot \delta(M) \cdot \mu(0)$ by appending some input action μ . The above result, together with the prefix-closure of $\delta_M(\text{nttraces}(q))$ and $\Delta_M(q)$, implies that these sets also coincide for such traces, i.e. the corollary holds. \square

This corollary means that if an implementation q can be assumed to be M -quiescent we may use the set of enriched observations $\delta_M(\text{nttraces}(q))$ to obtain $\Delta_M(q)$, whose definition is based on the unobservable timed transition system $\Delta(q)$. This will be the basis for our test derivation algorithm.

4 A real-time test generation framework

In this section we define the concept of real-time test cases, the nature of their execution, and the evaluation of their success or failure.

Definition 25.

- A test case t is a TLTS $\langle S, s_0, \text{Act}_\varepsilon \cup \{\delta\}, \rightarrow \rangle$ such that
 - t is deterministic and has bounded behaviour, i.e. $\exists N > 0$ such that for all computations $\sigma = \mu_1 \cdot \mu_2 \cdot \mu_3 \dots : |\{i \mid \mu_i \neq \varepsilon(d_i)\}| < \infty$ and $\Sigma\{d_i \mid \mu_i = \varepsilon(d_i)\} < N$
 - S contains the terminal states **pass** and **fail**, with $\text{init}(\text{pass}) = \text{init}(\text{fail}) = \emptyset$
 - for any state $t' \in S$ of the test case with $t' \neq \text{pass}, \text{fail}$, $\exists d > 0$ with $\text{init}(t' \text{ after } \varepsilon(d')) = \text{Act}_U \cup \{\varepsilon(e) \mid e = d - d'\}$ for all $d' < d$, $\text{init}(t' \text{ after } \varepsilon(d)) = \mu$ with $\mu \in \text{Act}_I$ or $\mu = \delta$
 - t does not have τ -transitions

The class of test cases over Act_I and Act_U is denoted as $\mathcal{TTEST}(Act_I, Act_U)$
 But we represent it similarly as a timed automata, only for simplify the notation

- A test suite \mathbf{T} is a set of test cases: $\mathbf{T} \subseteq \mathcal{TTEST}(Act_I, Act_U)$.

A test run of an implementation with a test case is modelled by the synchronous parallel execution of the test case with the implementation under test. This run continues until no more interactions are possible, i.e. until a deadlock occurs.

Definition 26. Let $t \in \mathcal{TTEST}(Act_I, Act_U)$ and $imp \in \text{TIOTS}(Act_I, Act_U)$ M -quiescent, then

- Running a test case t with an implementation imp is modelled by the parallel operator $\parallel : \mathcal{TTEST}(Act_I, Act_U) \times \text{TIOTS}(Act_I, Act_U) \rightarrow \text{TIOTS}(Act_I, Act_U)$ which is defined by the following inference rules:

$$\begin{array}{lcl} imp \xrightarrow{\tau} imp' & \vdash & t \parallel imp \xrightarrow{\tau} t \parallel imp' \\ t \xrightarrow{\delta} t' & \vdash & t \parallel imp \xrightarrow{\delta} t' \parallel imp \\ t \xrightarrow{\mu} t', imp \xrightarrow{\mu} imp', \mu \in Act & \vdash & t \parallel imp \xrightarrow{\mu} t' \parallel imp' \\ t \xrightarrow{\varepsilon(d)} t', imp \xrightarrow{\varepsilon(d)} imp' & \vdash & t \parallel imp \xrightarrow{\varepsilon(d)} t' \parallel imp' \end{array}$$

- A test run of t with imp , is a $\sigma \in \Delta_M$ of $t \parallel imp$ leading to a terminal state of t : σ is a test run of t and

$$imp \stackrel{def}{=} \exists imp' : (t \parallel imp \xrightarrow{\sigma} \mathbf{pass} \parallel imp') \text{ or } (t \parallel imp \xrightarrow{\sigma} \mathbf{fail} \parallel imp')$$

- An implementation imp **passes** test case t , if all their test runs lead to the **pass** state of t :

$$imp \text{ passes } t \stackrel{def}{=} \forall \sigma \in \Delta_M : \forall imp' : t \parallel imp \not\xrightarrow{\sigma} \mathbf{fail} \parallel imp'$$

- An implementation imp **passes** a test suite \mathbf{T} , if it **passes** all test cases in \mathbf{T} :

$$imp \text{ passes } \mathbf{T} \stackrel{def}{=} \forall t \in \mathbf{T} : imp \text{ passes } t$$

If imp does not pass the test suite, it fails if:

$$imp \text{ fails } \mathbf{T} \stackrel{def}{=} \exists t \in \mathbf{T} : imp \text{ passes } t.$$

Since an implementation can behave nondeterministically, different test runs of the same test case with the same implementation may lead to different terminal states and hence to different verdicts. An implementation **passes** a test case if and only if all possible test runs lead to the verdict **pass**.

4.1 Nondeterministic test case construction

For the description of test cases we use, as we already did before, a process-algebraic behaviour notation with a syntax inspired by LOTOS [ISO89]:

$$B \stackrel{def}{=} a; B \mid B + B \mid \Sigma B$$


where $a \in Act_\varepsilon$, \mathcal{B} is a countable set of behaviour expressions, and the axioms and the inference rules are:

$$\begin{array}{lcl}
a \in Act & \vdash & a; B \xrightarrow{a} B' \\
a = \varepsilon(d), d' < d & \vdash & \varepsilon(d); B \xrightarrow{\varepsilon(d')} \varepsilon(d - d'); B \\
a = \varepsilon(d) & \vdash & B \xrightarrow{\varepsilon(d)} B' \\
B_1 \xrightarrow{\mu} B'_1, \mu \in Act_\varepsilon & \vdash & B_1 + B_2 \xrightarrow{\mu} B'_1 \\
B_2 \xrightarrow{\mu} B'_2, \mu \in Act_\varepsilon & \vdash & B_1 + B_2 \xrightarrow{\mu} B'_2 \\
B \xrightarrow{\mu} B', B \in \mathcal{B}, \mu \in Act_\varepsilon & \vdash & \Sigma B \xrightarrow{\mu} B'
\end{array}$$

Moreover, we use $\mu(d)$ as syntactic sugar for $\varepsilon(d); \mu$.

Test case generation procedure We define a procedure to generate test cases from a given specification timed transition system. Similar to [J.T96] test cases result from the non-deterministic, recursive application of three test generation steps, corresponding to: (1) termination, (2) generation of an input, and (3) observation of output (including *quiescence*). It should be noted that the construction steps involve (negations of) predicates of the form $o(d) \in out_M(S)$, which on the general level of timed input-output transition systems are undecidable. The procedure given here, therefore, should be seen as a meta-algorithm that can be used to generate tests effectively for subclasses of TIOTS for which these predicates are decidable, such as timed automata [MS04,KMB03].


1. termination



The single state test case **pass** is always a sound test case. It stops the recursion in the algorithm, and thus terminates the test case.

2. inputs

choose $k \in (0, M)$ and $\mu \in Act_I$




$$\begin{aligned}
t := & \Sigma\{o_i(d_i); t_i \mid o_i \in Act_U \wedge o_i(d_i) \in out_M(S)\} \\
& + \mu(k); t_\mu \\
& + \Sigma\{o_j(d_j); \mathbf{fail} \mid o_j \in Act_U \wedge o_j(d_j) \notin out_M(S)\}
\end{aligned}$$

where x is a clock, k is a timed variable and t_i and t_μ are obtained by recursively applying the algorithm for $(S \text{ after } o_i(d_i))$ and $(S \text{ after } \mu(k))$, respectively.

Test case t is waiting for k time-units and treating to make and input (μ). If an output arrives from the implementation it checks; if it is an invalid response, i.e. $o_j(d_j) \notin out_M(S)$ then the test case terminates in **fail**; if it is a valid response after the timed pass then the test case continues recursively. If the time pass then the test makes the input (μ) and continues recursively.

3. waiting for outputs



$$\begin{aligned}
t := & \Sigma\{o_i(d_i); t_i \mid o_i \in Act_U \wedge o_i(d_i) \in out_M(S)\} \\
& + \Sigma\{\delta(M); t_\delta \mid \delta \in out_M(S \text{ after } \varepsilon(M))\} \\
& + \Sigma\{\delta(M); \mathbf{fail} \mid \delta \notin out_M(S \text{ after } \varepsilon(M))\} \\
& + \Sigma\{o_j(d_j); \mathbf{fail} \mid o_j \in Act_U \wedge o_j(d_j) \notin out_M(S)\}
\end{aligned}$$

where x is a clock and t_i and t_δ are obtained by recursively applying the algorithm for $(S \text{ after } o_i(d_i))$ and $(S \text{ after } \varepsilon(M))$, respectively.

Test case t is waiting for M time-units if an output arrive from the implementation it checks; if it is an invalid response, i.e. $o_j(d_j) \notin out_M(S)$ then the test case terminates in **fail**; if it is a valid response after the timed pass then the test case continues recursively. The observation of *quiescence* δ is treated separately, using the constant M given by the M -*quiescent* property.

Soundness The test generation procedure presented is sound with respect to the \mathbf{tioco}_M relation. This property is shown in the following theorem.

Theorem 27. *Let $spec \in TIOTS$, then for all M -quiescent $imp \in TIOTS$ and all test cases t obtained from $spec$ by the above procedure:*

$$imp \mathbf{tioco}_M spec \Rightarrow imp \text{ passes } t$$

Proof. Let imp be M -quiescent with $imp \mathbf{tioco}_M spec$, then we will show that for all $\sigma \in \Delta_M(spec)$ and all test cases t generated by the procedure from $spec$: if $(t||imp \xrightarrow{\sigma} t' || imp')$ then $(t' \neq \mathbf{fail})$. By induction on the length of σ :

- if $\sigma = \epsilon$ and $t||imp \xrightarrow{\epsilon} t' || imp'$
 - if t was constructed using case 1 in the first step, then $t||imp \xrightarrow{\epsilon} \mathbf{pass} || imp'$

- if t was constructed using cases 2 or 3 in the first step, then $t = t' \neq \mathbf{fail}$ and all derivations of $\xRightarrow{\xi}$ have the form:

$$t \parallel \mathit{imp} \xRightarrow{\xi} t \parallel \mathit{imp}'$$

- if $\sigma = \sigma' \cdot a$ and $t \parallel \mathit{imp} \xRightarrow{\sigma'} t'' \parallel \mathit{imp}'' \xRightarrow{a} t' \parallel \mathit{imp}' \wedge a = \mu(d)$, because t can do a there are only two possibilities to construct t' :
 - from case 2: $a = \mu(d) \wedge \mu \in \mathit{Act}_I \wedge d < M$, then $t \parallel \mathit{imp} \xRightarrow{\sigma} t' \parallel \mathit{imp}' \wedge t' \neq \mathbf{fail}$
 - from case 3: $a = \mu(d) \wedge a \in \mathit{out}_M(\mathit{imp} \mathbf{after} \sigma)$, then because $(\mathit{imp} \mathbf{tioco}_M \text{ spec})$: $\mu(d) \in \mathit{out}_M(\text{spec} \mathbf{after} \sigma)$, and thus $t \parallel \mathit{imp} \xRightarrow{\sigma} t' \parallel \mathit{imp}' \wedge t' \neq \mathbf{fail}$.

□

Exhaustiveness The test generation procedure is also exhaustive in the sense that for each non-conforming implementation a test case can be generated that detects the non-conformance.

Definition 28. Let $p \in \mathit{TIOTS}$, then

$$\sigma \in \Delta_M(p) \text{ is } \delta(M)\text{-saturated iff for all } \sigma' \text{ with } \sigma \delta_M \sigma' \text{ we have } \sigma = \sigma'.$$

Proposition 29. If σ is $\delta(M)$ -saturated, then for all $\mu(d) \neq \delta(M)$ occurring in σ we have $d < M$.

Proof. Directly from the definition of the δ_M -relation. □

Lemma 30. Let $\text{spec} \in \mathit{TIOTS}$, $\sigma \in \Delta_M(\text{spec})$ $\delta(M)$ -saturated, and t' a test case generated by the procedure for $(\text{spec} \mathbf{after} \sigma)$ then there exists a test case t generated from spec with $t \xRightarrow{\sigma} t'$.

Proof. By induction on the length of σ :

- $|\sigma| = 0$ then take $t = t'$
- suppose t exists for all σ with length n and let $\sigma = \sigma' \cdot a$ and $a = \mu(d)$
 - if $\mu \in \mathit{Act}_I$, using case 2 for the input μ :

$$t \xRightarrow{\sigma'} t'' \xRightarrow{a} t'$$
 - if $a \in \mathit{out}_M(\text{spec} \mathbf{after} \sigma)$, using case 3:

$$t \xRightarrow{\sigma'} t'' \xRightarrow{a} t'$$

□

Theorem 31. Let $\text{spec} \in \mathit{TIOTS}$, then for all M -quiescent $\mathit{imp} \in \mathit{TIOTS}$ with $\mathit{imp} \mathbf{tioco}_M \text{ spec}$ there exists a test case t generated from spec by the procedure such that:

$$\mathit{imp} \text{ passes } t.$$

Proof. If $\mathit{imp} \mathbf{tioco}_M \text{ spec}$ then there exists $\sigma \in \Delta_M(\text{spec}) : \mathit{out}_M(\mathit{imp} \mathbf{after} \sigma) \not\subseteq \mathit{out}_M(\text{spec} \mathbf{after} \sigma)$. Without loss of generality we can assume that σ is $\delta(M)$ -saturated. Then let $a \in \mathit{out}_M(\mathit{imp} \mathbf{after} \sigma) \setminus \mathit{out}_M(\text{spec} \mathbf{after} \sigma)$ and $\mathit{imp} \xrightarrow{\sigma} \mathit{imp}' \xrightarrow{a} \mathit{imp}''$. Let t' be the result of applying case 3 of the procedure to $(\text{spec} \mathbf{after} \sigma)$, and let t be the test case constructed out of t' and σ by the above lemma. Because $a \notin \mathit{out}_M(\text{spec} \mathbf{after} \sigma)$: $t \parallel \mathit{imp} \xrightarrow{\sigma \cdot a} \mathbf{fail} \parallel \mathit{imp}''$, imp passes t . □

The exhaustiveness of our test generation procedure is less useful than the corresponding result in the untimed case. There, it implies that the test generation algorithm, if repeatedly executed in a fair non-terminating manner, will generate all test cases in the limit, and therefore, in the limit, achieve full coverage with respect to **io** and the given specification *spec*.

Here, the number of potential test cases is uncountable because of the underlying continuous model of time, and no countable repetition of test generations suffices. It is possible, however, to obtain a version of the stronger form of exhaustiveness for real-time test generation as well by considering equivalence classes of (minimal) error traces. It can be shown that reasonable assumptions of our test generation procedure will hit each such equivalence class in the limit. This result will be reported in detail in a forthcoming publication.

5 Example

In the setting of timed automata, deciding the predicate $o_i(d_i) \in out_M(S)$ amounts to reachability analysis. For the simpler version of **tioco** based on timed trace inclusion (i.e. excluding *quiescence*) this has already been implemented in the tool environment IF [MS04], the UPPAAL-based testing tool TUPPAAL, and a real-time extension of TORX. We present an example of our test case generation based on a timed automaton model of a coffee machine, similar to the previous one, but with infinite behaviour due to cycles.

Example 32. Figure 3 shows two quirky coffee machines with time. The first one is a specification and the second one is a wrong implementation. To the right, there is a test case derived by the algorithm that can detect the error in the implementation. We suppose both machines are saturated with all input actions in each state. In the specification we show the δ -transitions, while in the implementation we detect them using $M = k$. We assume that $k > 1$.

The problem appears because:

$$out(\text{spec after } m?(1) \cdot c?(1) \cdot \delta(k) \cdot b?(1) \cdot c?(1)) = \{c![0, \infty]\}$$

and

$$out(\text{impl after } m?(1) \cdot c?(1) \cdot \delta(k) \cdot b?(1) \cdot c?(1)) = \{\delta(k)\}$$

where we use the notation $c![0, \infty]$ to denote that the output $c!$ can be at any time between 0 and ∞ .

6 Related work

As already indicated before this work is closely related to work carried out by Krichen et al. in [MS04], and closely related work by Larsen et al. [KMB03], who deal with a *quiescence*-free interpretation of timed **io** based on timed trace inclusion for timed automata. Our work shows how such results may be extended to deal with *quiescence*, and provides a general framework at the level of timed transition systems.

Previous attempts of extending testing with time include older work by Nielsen et al. in [BA01], for testing a subclass of timed automata called event-recording automata (ERA). The technique is based on the symbolic analysis of timed automata inspired by the UPPAAL model-checker, but lacks a suitable notion of implementation relation. Springintveld et al. in [JFP01] present an exhaustive testing method for deterministic timed automata with dense time, using the notion of a grid automaton that represents each clock region with a finite set

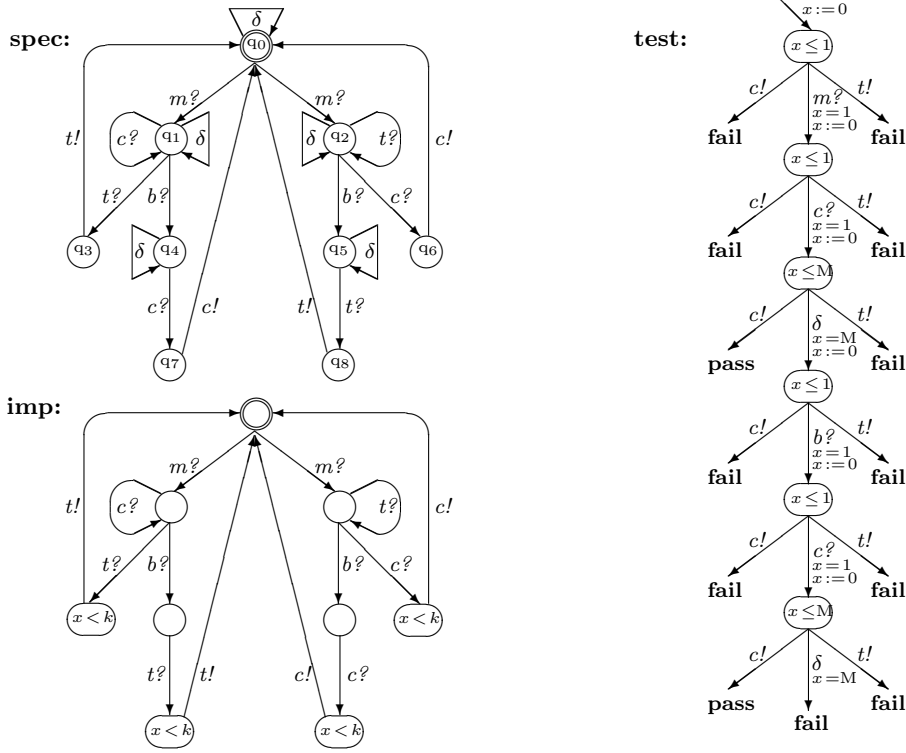


Fig. 3. A specification of a quirky coffee machine with time, an implementation with $M = k$, and a test case derived from the specification.

of clock valuations. Although being exact, the grid method is impractical because it generates “an astronomically large number of test sequences” [JFP01]. Cardell-Oliver presents a method for networks of deterministic timed automata extended with integer data variables [RO02], where only a part of the system is visibly using test views, so that a test is never exhaustive.

Several authors have tried to obtain good specification coverage for their test methods by adapting transition-tour methods from classical FSM-based testing [ENRF98, TAKR99].

Clarke and Lee [DI97] use the algebra of communicating shared resources (ACSR) on a discrete time base. ACSR allows non-deterministic specifications, the use of internal events and priorities. For testing, only boundary points of the time domain are selected. Cleaveland et al. propose a testing method for probabilistic processes on a discrete model of time [RIPS96] that bears a close resemblance to the classical testing theory of Hennessy and De Nicola [RM83]. Mandrioli et al. use temporal logic with arithmetic on a discrete time base [DSA95].

7 Conclusion and future work

In this paper we have presented an extension of Tretmans’ **ioco** theory and algorithm for test generation for input-output transition systems to real-time systems. Our treatment is based on an operational interpretation of the notion of *quiescence* that gives rise to a family of implementation relations parameterized by observation durations M for *quiescence*. These relations detect differences in behaviour after the execution of suspension traces provided that the observations of *quiescence* all take longer than the stipulated duration M , but may not detect differences in refusal behaviour that require shorter observations of *quiescence*.

It is shown how this theory may be used to test real-time implementations under the assumption that the absence of system interaction with its environment for M time units implies *quiescence*. We have defined a nondeterministic (M -parameterized) test generation framework that generates test cases that are sound with respect to the corresponding implementation relation \mathbf{tioco}_M . The test generation is also exhaustive in the sense that for each non-conforming implementation a test case can be generated that can detect the non-conformance.

The framework can be effectively instantiated for subclasses of timed input-output transition systems for which $out_M(spec \text{ after } \sigma)$ is computable, as is the case for timed automata. Using standard symbolic state space representation in the form of difference bounded matrices [D.D90], a real-time version of TORX for timed automata models is being implemented.

The work presented here can be extended in a number of ways. As already indicated, it is possible to show a stronger exhaustiveness result for the test generation procedure based on an appropriate notion of equivalence of error traces. The generation procedure will hit each such class in the limit, provided that the error class is not negligible, i.e. it must have positive measure in some appropriate sense.

Another extension is to relax the requirement that there must be a uniform observation deadline M for *quiescence*. Obvious alternatives that we are studying are:

- the observation parameter $M(\sigma)$ is a function of the behaviour (trace) σ observed so far. This allows us to model sequential phases of *quiescence*, i.e. slow vs. quick response times.
- the observation parameter $M(C_i)$ is a function of the communication channel C_i on which output is being observed. This allows us to have model different kinds of response times for different communication channels with the system under test, and would correspond to a real-time extension of the \mathbf{mioco} implementation relation of [ELJ98].

Our real-time theory inherits its focus on control aspects of system behaviour from the existing \mathbf{ioco} theory. Ultimately, it will be important to combine this testing theory with methods for testing the static data aspects of systems. It will be interesting to see to what extent the symbolic representation of data types can be combined with symbolic representations of time.

In a more general vein, one can say that the development of a real-time testing theory forces us to confront modelling issues with respect to physical aspects of time and implementation. From a physical point of view, for example, it is questionable whether negligible behaviour can be implemented. This has also implications for specification formalisms that can be used to specify such behaviour, e.g. timed automata can define negligible behaviour by using guards that force behaviour to go through specific points in time, such as $x = 3$. It would seem that realistic specifications and/or implementation relations allow for tolerances in the evaluation of clock conditions. This would then introduce a third source of non-determinism in the testing theory of real-time systems. At any rate, a more systematic study of the formal aspects of tolerance and robustness is definitely needed.

References

- [AJR⁺99] A.Belifante, J.Feenstra, R.deVries, J.Tretmans, N.Goga, L.Feijs, S.Mauw, and L.Heerink. Formal test automation: A simple experiment. In G.Csopaki, S.Dibuz, and K.Tarnay, editors, *Int. Workshop on Testing of Communicating Systems 12*, pages 179–196. Kluwer, 1999.
- [BA01] B.Nielsen and A.Skou. *Automated Test Generation from Timed Automata*. TACAS 2001: 343-357, 2001.
- [BBB]
- [D.D90] D.Dill. Timing assumptions and verification of finite-state concurrent systems. In *Proceedings of the int. workshop on Automatic verification methods for finite state systems*, pages 197–212. Springer-Verlag NY, Inc., 1990.
- [DI97] D.Clarke and I.Lee. Automatic test generation for the analysis of a real-time system: Case study. In *IEEE Real Time Technology and Applications Symp.*, pages 112–124, 1997.
- [DSA95] D.Mandrioli, S.Morasca, and A.Morzenti. Generating test cases for real-time systems from logic specifications. *TOMS*, 13(4):365–398, Nov. 1995.
- [ELJ98] E.Brinksma, L.Heerink, and J.Tretmans. Factorized test generation for multi input/output transition systems. In A.Petrenko and N. Yevtushenko, editors, *Int. Workshop on Testing of Communicating Systems 11*, pages 67–82. Kluwer, 1998.
- [ENRF98] En-NouaaryA, R.Dssouli, and F.Khendek. Timed test cases generation based on state characterization technique. In *19th IEEE Real-Time Systems Symp.*, pages 220–229, 1998.
- [ISO89] ISO8807. *Information processing systems, Open Systems Interconnection, LOTOS, A formal description technique based on the temporal ordering of observational behaviour*. Int. Organization for Standardization, 1989.
- [JCCTC96] J-C.Fernandez, C.Jard, T.Jeron, and C.Viho. Using on-the-fly verification techniques for the generation of test suites. In *Copmputer Aided Verification CAV’96. LNCS 1102, Springer-Verlan*. R.Alur and T.A.Hezinger, 1996.
- [JCCTC97] J-C.Fernandez, C.Jard, T.Jeron, and C.Viho. An experiment in automatic generation of test suites for protocols with verification technology. In *Science of Computer Programming - Special Issue on COST247, Verification and Validation Methods for Formal Descriptions, 29(1-2)*, pages 123–146, 1997.
- [JE03] J.Tretmans and E.Brinksma. Torx: Automated model-based testing. In *First European Conference on Model-Driven Software Engineering, Nuremberg*. A.Hartmann and K.Dussa-Ziegler, 2003.
- [JFP01] J.Springintveld, F.Vaandrager, and P.D’Argenio. Testing timed automata. *Theoretical Computer Science*, 254(1–2):225–257, 2001.
- [J.T96] J.Tretmans. Test generation with inputs, outputs and repetitive quiescence. In *Software-Concepts and Tools, 17(3)*, pages 103–120. Also: Technical Report N0. 96-26, Center for Telematics and Information Technology, University of Twente, The Netherlands, 1996.
- [KMB03] K.Larsen, M.Mikucionis, and B.Nielsen. Real-time system testing on-the-fly. In K.Sere, M.Walden, and A.Karlsson, editors, *The 15th Nordic Workshop on Programming Theory (NWPT)*, Åbo Akademi University, Turku, Finland, oct 2003. Extended abstract.
- [MS04] M.Krichen and S.Tripakis. Black-box conformance testing for real-time systems. In *SPIN 2004*, pages 109–126. Springer-Verlag Heidelberg, 2004.
- [RIPS96] R.Cleaveland, I.Lee, P.Lewis, and S.Smolka. A theory of testing for soft real-time processes, 1996.
- [R.J93] R.J.vanGlabbeek. The linear time-branching time spectrum ii (the semantics of sequential systems with silent moves). In *CONCUR’93. LNCS 715*, pages 66–81. E.Best, 1993.
- [R.L90] R.Langerak. A testing theory for lotos using deadlock detection. In *Proceedings of the IFIP WG 6.1 Ninth int. Symp. on Protocol Spec., Testing, and Verification*, pages 87–98. IFIP, 1990.
- [RM83] R.deNicola and M.Hennessy. Testing equivalences for processes. In *ICALP83*, volume 154, 1983.
- [RO02] R.Cardell-Oliver. Conformance test experiments for distributed real-time systems. In *Proceedings of the int. symp. on Software testing and analysis*, pages 159–163. ACM Press, 2002.
- [TAKR99] T.Higashino, A.Nakata, K.Taniguchi, and R.Cavalli. Generating test cases for a timed i/o automaton model. In *IWTCS 1999*, pages 197–214, 1999.