

A Tool for the Syntactic Detection of Zeno-timelocks in Timed Automata

Howard Bowman¹, Rodolfo Gomez^{2,3} and Li Su⁴

*Computing Laboratory
University of Kent,
Canterbury, United Kingdom*

Abstract

Timed automata are a very successful notation for specifying and verifying real-time systems, but timelocks can freely arise. These are counter-intuitive situations in which a specifier's description of a component automaton can inadvertently prevent time from passing beyond a certain point, possibly making the entire system stop. In particular, a zeno-timelock represents a situation where infinite computation is performed in a finite period of time. Zeno-timelocks are very hard to detect for real-time model checkers, e.g. UPPAAL and Kronos. We have developed a tool which can take an UPPAAL model as input and return a number of loops which can potentially cause zeno-timelocks. This tool implements an algorithm which refines a static verification approach introduced by Tripakis. We illustrate the use of this tool on a real-life case-study, the CSMA/CD protocol.

Keywords: Timed Automata, Zeno-timelocks, UPPAAL.

1 Introduction

Timed automata are one of the success stories of formal methods. This is largely because they have proved to be amenable to automatic verification using symbolic model checking. Perhaps the most prominent incarnation of such symbolic model checking is the UPPAAL tool [7], which has been used

¹ Email: H.Bowman@kent.ac.uk

² Supported by the ORS Award Scheme.

³ Email: rsg2@kent.ac.uk

⁴ Email: ls68@kent.ac.uk

to verify a number of non-trivial real-time systems (visit www.uppaal.com for examples and documentation). Despite these successful applications of timed automata model checking, there are some difficulties with the approach. Perhaps the most important is that timelocks can freely arise and furthermore, it can be very difficult to determine that a non-trivial system is free from such timelocks.

Informally speaking a system can timelock if a state can be reached where no possible subsequent run allows time to diverge, i.e. pass by an infinite amount. It is important to note that timelocks can arise for a number of reasons and that different classes of timelock need to be handled in different ways. In particular, we can distinguish between the following two classes of timelock.

Time-actionlocks are states in which neither time or action transitions can be performed, which typically arise when there is a conflict between the urgency and the synchronisation properties of the system, i.e. when a location invariant ensures that a component automaton must perform a half-action at a time at which no other component automaton is offering a matching half-action. Thus, the synchronisation *must* happen (due to the urgency constraint) at a point at which it is not enabled.

Zeno-timelocks are situations in which time is unable to pass beyond a certain point, but actions continue to be performed. Thus, the system is continuing to evolve but none of these evolutions will enable time to diverge. The hallmark of such paradoxical runs is that an infinite number of actions are performed in a finite period of time.

It is also important to realise that timelocks are quite different from actionlocks, which are the analogue of deadlocks in untimed specifications. Critically, actionlocks allow time to pass; the automaton may not be able to perform any further “useful” computation, but it can still pass time, which means that it does not prevent other component automata from passing time. The fact that local actionlocks do not propagate globally is the reason why actionlocks are much more palatable than timelocks. Global propagation in the timelock situation arises because global time passing is dependent upon local time passing. As illustrated by the fact that a collection of timed automata can only pass time by t time units if all component automata can pass time by t time units. Thus, effectively, automata synchronise on the passage of time.

In previous papers we have considered the timelock problem, classified different types of timelocks and highlighted solutions corresponding to the needs of these different classes [3,2,4]. Our main interest here though is with zeno-timelocks; we present an analytical method to ensure absence of zeno-timelocks which builds upon the notion of *strong non-zenoness* introduced by Tripakis

[9]. We show how Tripakis’ results can be extended, broadening the class of timed automata specifications which can be guaranteed to be free from zeno-timedlocks. In particular, the relationship between strong non-zenoness and synchronising components is analysed in more detail. Moreover, we present a tool that we have developed which implements this syntactic verification on UPPAAL-like timed automata specifications. Also, new syntactic properties, in the spirit of strong non-zenoness, are presented which also ensure zeno-timedlock freedom. This sufficient-only approach can only guarantee that zeno-timedlocks *do not* occur, but it presents two important advantages: a) it works at a syntactic level, and thus it is more efficient than reachability analysis, and b) it identifies all potential sources of zeno-timedlocks directly on the timed automata models. Therefore, even if the method fails to recognise that a model is free from zeno-timedlocks, it helps the user in narrowing the analysis to specific parts of the model.

To the best of our knowledge, no other tool implements syntactic checks for zeno-timedlock freedom. For example, UPPAAL does not support any form of zeno-timedlock checking. Some other tools do better, e.g. Kronos [10], but they suffer from other problems. For example, Kronos is not as usable a tool as UPPAAL. UPPAAL presents a well-developed GUI, a rich modelling language, a graphical simulator, and a fast verifier, among other features. Kronos can verify the *TCTL* formula $\forall \square \exists \diamond_{=1} true$, whose satisfaction represents a sufficient and necessary condition to ensure timedlock freedom, but its verification is based on reachability analysis. Thus, it could be that for some specifications, checking timedlock freedom in Kronos would be the most expensive requirement to check and the need to check it could prevent a complete verification.

We begin by defining timed automata and their syntax (see Section 2). Then we present a summary of our classification of deadlocks, and particularly timedlocks (see Section 3). Following this we highlight the theory behind our zeno-timedlock checking approach (see Section 4). Then the main body of the paper describes our tool and presents a case study of zeno-timedlock checking via the verification of a CSMA/CD protocol (Section 5). Finally, Section 6 presents concluding remarks, and an appendix is provided which includes proofs of the theorems presented in the paper. An extended version of this paper appears in [5].

2 Timed Automata Notation

Basic Sets

CA is a set of *completed* (or internal) actions. $HA = \{ a?, a! \mid a \in CA \}$ is a

set of *half* (or *uncompleted*) actions. These give a simple CCS style [6] point-to-point communication similar, for example, to the synchronisation primitives found in UPPAAL [7]. Thus, two actions, $a?$ and $a!$ can synchronise and generate a completed action a . $\mathbb{A} = HA \cup CA$ is the set of *all* actions. \mathbb{R}^+ denotes the positive reals without zero and $\mathbb{R}^{+0} = \mathbb{R}^+ \cup \{0\}$. \mathbb{C} is the set of all clock variables, which take values in \mathbb{R}^{+0} . CC is a set of clock constraints of the form $x \sim n$, $x - y \sim n$ or $\phi_1 \wedge \phi_2$, where $n \in \mathbb{N}$, $x, y \in \mathbb{C}$, $\phi_1, \phi_2 \in CC$ and $\sim \in \{<, >, =, \leq, \geq\}$. Also if $C \subseteq \mathbb{C}$ we write CC_C for the set of clock constraints generated from clocks in C . $\mathbb{V} = \mathbb{C} \rightarrow \mathbb{R}^{+0}$ is the space of possible clock valuations and $\mathbb{V}_C = C \rightarrow \mathbb{R}^{+0}$ is the space of clock valuations for clocks in C . \mathbb{L} is the set of all possible automata locations.

Timed Automata

An arbitrary element of \mathcal{A} , the set of all timed automata, has the form (L, l_0, T, I, C) , where the elements are as follows. $L \subseteq \mathbb{L}$ is a finite set of locations; $l_0 \in L$ is a designated *start location*. C is the set of clocks of the timed automaton. $T \subseteq L \times \mathbb{A} \times CC_C \times \mathbb{P}(C) \times L$ is a transition relation (where $\mathbb{P}(S)$ denotes the powerset of S). A typical element of T would be, (l_1, a, g, r, l_2) , where $l_1, l_2 \in L$ are automaton locations; $a \in \mathbb{A}$ labels the transition; $g \in CC_C$ is a guard; and $r \in \mathbb{P}(C)$ is a reset set. $(l_1, a, g, r, l_2) \in T$ is typically written, $l_1 \xrightarrow{a, g, r} l_2$, stating that the automaton can evolve from location l_1 to l_2 if the (clock) guard g holds and in the process action a will be performed and all the clocks in r will be set to zero. $I : L \rightarrow CC_C$ is a function which associates an invariant with every location. Informally, the automaton can remain on a given location only as long as the invariant is true. Thus, invariants are used to model urgency: (enabled) outgoing transitions must be taken immediately when the corresponding location invariant is false. We will be precise about the interpretation of invariants when we discuss the semantics of TAs shortly, however, it is important to understand the difference between the role of guards and invariants. In this respect we can distinguish between *may* and *must* timing. Guards express *may* behaviour, i.e. they state that a transition is possible or in other words *may* be taken. However, guards cannot “force” transitions to be taken. In contrast, invariants define *must* behaviour. This *must* aspect corresponds to *urgency*, since an alternative expression is that when an invariant expires, outgoing transitions must be taken straightaway.

Semantics

Timed automata are semantically interpreted over transition systems which are triples, (S, s_0, \Rightarrow) , where $S \subseteq \mathbb{L} \times \mathbb{V}$ is a set of states; $s_0 \in S$ is a start state;

and $\Rightarrow \subseteq S \times Lab \times S$ is a transition relation, where $Lab = \mathbb{A} \cup \mathbb{R}^+$. Thus, transitions can be of one of two types: *discrete transitions*, e.g. (s_1, a, s_2) , where $a \in \mathbb{A}$ and *time transitions*, e.g. (s_1, d, s_2) , where $d \in \mathbb{R}^+$ and the passage of d time units is denoted. Transitions are written: $s_1 \xrightarrow{a} s_2$ respectively $s_1 \xrightarrow{d} s_2$.

For a clock valuation $v \in \mathbb{V}_C$ and a delay d , $v+d$ is the clock valuation such that $(v+d)(c) = v(c) + d$ for all $c \in C$. For a reset set r , we use $r(v)$ to denote the clock valuation v' such that $v'(c) = 0$ whenever $c \in r$ and $v'(c) = v(c)$ otherwise. v_0 is the clock valuation that assigns all clocks to the value zero.

The semantics of a timed automaton $A = (L, l_0, T, I, C)$ is a transition system, (S, s_0, \Rightarrow) , where $S = \{s' \in L \times \mathbb{V}_C \mid \exists s \in S, y \in Lab. s \xrightarrow{y} s'\} \cup \{[l_0, v_0]\}$ is the set of reachable states, $s_0 = [l_0, v_0]$ and \Rightarrow is defined by two inference rules ($I(l_0)(v_0)$ is required to hold):

$$\frac{l \xrightarrow{a,g,r} l' \quad g(v) \quad I(l')(r(v))}{[l, v] \xrightarrow{a} [l', r(v)]} \qquad \frac{\forall d' \leq d. I(l)(v + d')}{[l, v] \xrightarrow{d} [l, v + d]}$$

The first rule gives an interpretation to invariants such that locations cannot be entered if the corresponding invariant is false. This interpretation, usually known as the *strong-invariant interpretation*, is the one adopted by UPPAAL and also assumed by our non-urgency properties presented in Section 4.

Parallel Composition

We assume our system is described as a network of timed automata. These are modelled by a vector of automata⁵ denoted, $|A = |\langle A[1], \dots, A[n] \rangle$ where $A[i]$ is a timed automaton. In addition, we let u, u' , etc, range over the set \mathbb{U} of vectors of locations, which are written, $\langle u[1], \dots, u[n] \rangle$. We use a substitution notation as follows: $\langle u[1], \dots, u[j], \dots, u[n] \rangle [u[j]'/u[j]] = \langle u[1], \dots, u[j-1], u[j]', u[j+1], \dots, u[n] \rangle$ and we write $[u[j]'/u[j]]$ as $[j'/j]$ and $u[i'_1/i_1] \dots [i'_m/i_m]$ as $[i'_1/i_1, \dots, i'_m/i_m]$.

If $\forall i(1 \leq i \leq n). A[i] = (L_i, l_{i,0}, T_i, I_i, C_i)$ then the product automaton, which characterises the behaviour of $|\langle A[1], \dots, A[n] \rangle$ is given by, (L, l_0, T, I, C) where $L = \{ |u \mid u \in L_1 \times \dots \times L_n \}$, $l_0 = |\langle l_{1,0}, \dots, l_{n,0} \rangle$, T is as defined by the following two inference rules, $I(|\langle u[1], \dots, u[n] \rangle) = I_1(u[1]) \wedge \dots \wedge I_n(u[n])$ and $C = C_1 \cup \dots \cup C_n$.

$$\frac{u[i] \xrightarrow{x^?, g_i, r_i} u[i]' \quad u[j] \xrightarrow{x!, g_j, r_j} u[j]'}{|u \xrightarrow{x, g_i \wedge g_j, r_i \cup r_j} |u[i]'/i, j'/j|} \qquad \frac{u[i] \xrightarrow{x, g, r} u[i]' \quad x \in CA}{|u \xrightarrow{x, g, r} |u[i]'/i|}$$

⁵ Although our notation is slightly different, our networks can be related, say, to the process networks used in UPPAAL.

where $1 \leq i \neq j \leq |u|$. Note, we write $x \leq k \neq r \leq y$ in place of $x \leq k \leq y \wedge x \leq r \leq y \wedge k \neq r$.

3 Classification of Deadlocks

In a very broad sense, deadlocks are states where the system is unable to progress further. We would expect the system to be able to run forever hence deadlocks can be seen as error situations. In untimed systems, deadlocks are states where the system will never be able to perform an action. However, in timed automata, the range of transitions has been broadened to time passing and discrete transitions (actions). Consequently, in this setting the ways of violating the requirements of progress can vary. So deadlocks in timed automata can be of different types. We will highlight these different types and in addition, as an assessment of the state of the art we will also consider the means that UPPAAL provides for checking for such locks.

Before giving the formal definitions of various types of deadlocks, we briefly review the terminology we will use, this is largely inherited from [3,9]. Given $s = [l, v]$, we will write $s+d$ instead of $[l, v+d]$. Also, we will write $s \xrightarrow{x}$ to denote $\exists s'. s \xrightarrow{x} s'$. A *run* of $A \in TA$ starting from state s_0 is a finite or infinite sequence: $\rho = s_0 \xrightarrow{d_0} s_0 + d_0 \xrightarrow{a_1} s_1 \dots s_{n-1} \xrightarrow{d_{n-1}} s_{n-1} + d_{n-1} \xrightarrow{a_n} s_n \xrightarrow{d_n} \dots$, where $\forall 0 \leq i \leq n. s_i \in S, a_i \in \mathbb{A}; \forall 0 \leq i \leq n-1. d_i \in \mathbb{R}^{+0}; d_n \in \mathbb{R}^{+0} \cup \{\infty\}$ ($s_n \xrightarrow{\infty}$ denotes $\forall t \in \mathbb{R}^{+0}. s_n \xrightarrow{t}$). Notice that infinite runs contain an infinite number of discrete transitions (i.e. actions). Let $Tr(A)$ denote the set of all runs of A , and define the function $delay(\rho)$, $\rho \in Tr(A)$ as the sum of all delays d_i in ρ . If $delay(\rho) = \infty$ we say that the ρ is a *divergent run*.

Generally speaking, *actionlocks* are states where no discrete transition can be performed, while *timelocks* are states where time cannot pass beyond a certain point. Formally, given $A \in TA$, a state $s = [l, v]$ is an actionlock if $\forall d \in \mathbb{R}^{+0}. [l, v+d] \in S \Rightarrow \nexists a \in \mathbb{A}. [l, v+d] \xrightarrow{a}$. Thus, however long the system idles in location l no action can be performed. However, a state s is a timelock if there is no divergent run $\rho \in Tr(A)$ starting at s . A timed automaton A is actionlock-free (timelock-free) if none of its reachable states is an actionlock (timelock). Actionlocks and timelocks can be further refined as *pure-actionlocks*, *time-actionlocks* or *zeno-timelocks* (or pure timelocks), which are explained next.

Definition 3.1 (Pure-actionlock) Pure-actionlocks are states of a system where it cannot perform any discrete transitions, but can still pass time arbitrarily. Given $A \in TA$, a state $s = [l, v]$ is a *pure-actionlock* if $\forall d \in \mathbb{R}^{+0}. [l, v+d] \in S \wedge \nexists a \in \mathbb{A}. [l, v+d] \xrightarrow{a}$.

Fig. 1a shows an example of a timed automaton with a pure actionlock: no action is enabled once the automaton reaches location S0, however time is not prevented from passing.

Definition 3.2 (Time-actionlock) Time-actionlocks are states where neither discrete nor time transitions can be performed. Given $A \in TA$, a state s is a *time-actionlock* if $\nexists a \in \mathbb{A}, d \in \mathbb{R}^+. s \xrightarrow{a} \vee s \xrightarrow{d}$.

An example of a time-actionlock is shown in Fig. 1b. The upper automaton must perform an action $a!$ before more than 5 time units have passed, while the bottom one can only perform an $a?$ after 5 units have passed. The system, then, enters in a time-actionlock immediately after 5 time units have elapsed.

Definition 3.3 (Zeno-timelock) Given $A \in TA$, a *zeno* run is an infinite run $\rho \in Tr(A)$ s.t. $delay(\rho) \neq \infty$. A state s is a *zeno-timelock* if a) there is at least one infinite run starting at s , b) all infinite runs starting at s are zeno, and c) there is no run $\rho' \in Tr(A)$ starting at s s.t. $delay(\rho') = \infty$ ⁶.

When a system enters a zeno-timelock, transitions can still be performed (which can be either discrete or time transitions) but time cannot pass beyond a certain point. This models a situation where the system performs an infinite number of transitions in a finite period of time. Fig. 1c shows a zeno-timelock, where transition a must be performed an infinite number of times before more than 5 time units have passed.

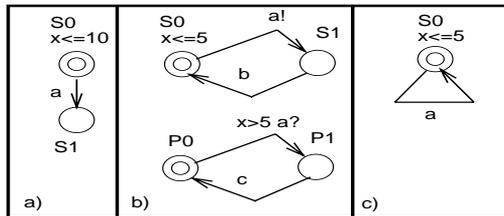


Fig. 1. Classification of deadlocks

Discussion

One reason for presenting our classification is that we believe that different types of deadlocks bring different types of problems and, hence, should be treated differently. Firstly, although pure-actionlocks may be undesirable within the context of a particular specification, they are not of themselves counter-intuitive situations. It is wholly reasonable that a component or a system might reach a state from which it cannot perform any actions, as long as

⁶ According to our definition of run, ρ' is not necessarily infinite.

such an actionlock does not stop time. Thus, although analytical tools which detect pure-actionlocks certainly have value, we do not believe there is any fundamental reason why actionlocks should be prevented (by construction) at the level of the specification notation. In contrast, we are strongly of the opinion that time-actionlocks are counter-intuitive. In particular, and as previously discussed, a local “error” in one component has a global effect on the entire system, even if the remainder of the system has no actions in common with the timelocked component. Because of these particularly counter-intuitive aspects, we believe that time-actionlocks should be prevented *by construction*, i.e. the timed automata model should be reinterpreted in such a way that time-actionlocks just cannot arise. Bowman [3] presents such a method for Timed Automata with Deadlines [1]. Finally, to come to zeno-timelocks. Our position here is that analytical methods should be provided to check on a specification by specification basis whether zeno-timelocks occur. Our reasons for advocating this approach are largely pragmatic, since it is not clear how the timed automata model could be changed in order to constructively prevent such situations. In particular, any mechanism that ensured at the level of the semantics that a minimum time (say ϵ) was passed on every cycle, would impose rigid constraints on the specifiers ability to describe systems abstractly⁷. Section 4 considers just such an analytical method for detecting zeno-timelocks.

4 Zeno-timelocks

We now present an analytical method to ensure absence of zeno-timelocks which builds upon the notion of *strong non-zenoness* introduced by Tripakis [9]. We show how Tripakis’ results can be extended to guarantee zeno-timelock freedom for systems which may not be *strongly non-zeno*. In particular, the relationship between strong non-zenoness and synchronising components is analysed in more detail. Also, new syntactic properties, in the spirit of strong non-zenoness, are presented which also ensure zeno-timelock freedom.

The strong non-zenoness property, which we recall below, represents a sufficient but not necessary condition to ensure zeno-timelock freedom; systems which are strongly non-zeno are guaranteed to be free from zeno-timelocks, but there exists some systems which are free from zeno-timelocks but are not strongly non-zeno.

Definition 4.1 (Strong non-zenoness) Given $A \in TA$, a *structural loop* in A is a sequence of locations and edges in A , $l_0 \xrightarrow{a_1, g_1, r_1} l_1 \xrightarrow{a_2, g_2, r_2} \dots \xrightarrow{a_n, g_n, r_n} l_n$,

⁷ Note that early versions of timed CSP did employ exactly such an approach.

s.t. $l_0 = l_n$. A is *strongly non-zero* if for every such loop there exists a clock $c \in C$, $\epsilon \in \mathbb{R}^+$ and $0 \leq i, j \leq n$ s.t. (1) $c \in r_i$ and (2) c is bounded from below in step j , i.e. $g_j \Rightarrow c > \epsilon$. Every loop which satisfies these properties is also called strongly non-zero. Strong non-zenoness guarantees absence of zeno-timedlocks, and it is preserved by parallel composition. Lemma 4.2 below formalises these results [9].

Lemma 4.2 *If $A \in TA$ is strongly non-zero then $Tr(A)$ does not contain zeno-timedlocks. Moreover, if $A[1], \dots, A[n] \in TA$ are strongly non-zero then $|A$ is also strongly non-zero.*

Lemma 4.2 suggests a static verification method; a system is free from zeno-timedlocks if all its components are strongly non-zero or in other words, if every loop in every component is strongly non-zero. This result is justified by the structure of the product automaton, where every loop is the result of either two loops with matching half actions in two different component automata (we refer to them as *synchronising loops*), or a loop with no half actions in a component automaton (called *complete loops*). Since a) every component loop is strongly non-zero, b) strong non-zenoness depends only on the existence of a clock which is bounded from below in a given guard in the loop, and also reset at some point in the loop, and c) these conditions are preserved in the edges of resulting loops in the product automaton, then every loop in the product automaton is strongly non-zero. But notice that a strongly non-zero loop in a component is, in fact, “preserved” in every loop in the product which results from the synchronisation of this loop with any other loop in a different component, whether this is also strongly non-zero or not. Therefore, synchronisation between a strongly non-zero loop and *any* other loop must also be considered “safe”. This may have a considerable impact from the user’s side: a system will no longer be considered unsafe just because there is a loop in one of its components which is not strongly non-zero (this happens if we analyse the system according to Lemma 4.2). Instead, we can pair all synchronising loops in the collection of components, and for each pair, ask just for one loop to be strongly non-zero. It is also required that all loops which do not contain half-actions are strongly non-zero, because these loops are preserved in the product automaton, but the benefits of this approach are still evident. We have found, then, that the requirements imposed by Lemma 4.2 to ensure absence of zeno-timedlocks can be “weakened” to consider just a subset of all structural loops appearing in the component automata. This result is formalised by Lemma 4.3 below (proof is given in the appendix).

Lemma 4.3 *Let HL be the set of all pairs of synchronising loops in*

$A[1], \dots, A[n]$, and, respectively, CL the set of all complete loops. If (at least) one loop in every pair of HL is strongly non-zeno and all loops in CL are strongly non-zeno then $|A$ is also strongly non-zeno and thus free from zeno-timelocks.

We now present two other properties, *location non-urgency* and *reset non-urgency* which also work at the level of the timed automata syntax, and represent sufficient-only conditions. However, they can guarantee that a system is free from zeno-timelocks even when it may not be strongly non-zeno; in this way the scope of syntactic detection of zeno-timelock free systems is further broadened. The intuition is the same as that which underlies strong non-zenoness: we have to show for any state s that if there exist some infinite runs starting at s , then at least one of them must diverge (therefore s is not a zeno-timelock). Now by definition, infinite runs must necessarily traverse some loop an infinite number of times (otherwise the run could not contain an infinite number of discrete transitions). Therefore, we just need to ensure that time can pass by at least $\epsilon \in \mathbb{R}^+$ time units on every iteration of any loop (where ϵ is considered a constant value). Because these conditions focus on invariants, they cover some kinds of safe loops which are not strongly non-zeno. In the following definitions, we use $Clocks(I(l)) \subseteq C$ to denote the set of clocks appearing in the invariant expression $I(l)$ (where l is a location).

Definition 4.4 (Location non-urgency) $A \in TA$ is called *location non-urgent* if in every structural loop there is a location where either the invariant is *True* or every clock appearing in the invariant has no upper bound. For example, *True* and $x > 1$ (where $x \in C$) can be two such invariants. Formally, let $l_0 \xrightarrow{a_1, g_1, r_1} l_1 \xrightarrow{a_2, g_2, r_2} \dots \xrightarrow{a_n, g_n, r_n} l_n$, s.t. $l_0 = l_n$ be a structural loop in A . A is called *location non-urgent* if for every such structural loop there exists $0 \leq i \leq n$ s.t. $\exists d \in \mathbb{R}^+ . \forall c \in Clocks(I(l_i)), v \in \mathbb{V}_C. (v(c) > d \Rightarrow I(l_i)(v))$ (notice that this formula vacuously holds for *True* invariants, since $Clocks(True) = \emptyset$). These loops are also called location non-urgent.

Definition 4.5 (Reset non-urgency) $A \in TA$ is called *reset non-urgent* if in every structural loop there is a location where at least one clock in the invariant has a non-zero lower bound, and this clock is reset in the loop. Formally, let $l_0 \xrightarrow{a_1, g_1, r_1} l_1 \xrightarrow{a_2, g_2, r_2} \dots \xrightarrow{a_n, g_n, r_n} l_n$ s.t. $l_0 = l_n$, be a structural loop in A . A is called *reset non-urgent* if for every such structural loop there exists $0 \leq i \leq n$ s.t. $\exists d \in \mathbb{R}^+, c \in Clocks(I(l_i)), v \in \mathbb{V}_C. (v(c) = d \wedge I(l_i)(v) \wedge (\forall v' \in \mathbb{V}_C. v'(c) < d \Rightarrow \neg I(l_i)(v')) \wedge \exists j (0 \leq j \leq n). c \in r_j)$. These loops are also called reset non-urgent.

The following lemma states the relation between location non-urgency, reset non-urgency and zeno-timelock freedom (proof is given in the appendix).

Lemma 4.6 *If $A \in TA$ is either location non-urgent or reset non-urgent, then it is also free from zeno-timelocks.*

Location non-urgency is not compositional, i.e. the product of location non-urgent automata is not guaranteed to be free from zeno-timelocks; but we believe this property would be of use when applied to the product automaton. Its benefits are even more evident when we consider that its application to the product automaton may be less “expensive” than a semantic-based check [9]. On the other hand reset non-urgency is compositional, but it has not been applicable to the specifications we have been working with. Nevertheless, it remains an interesting alternative given the fact that (at least in principle) invariants with lower-bounds (e.g. $1 < x \leq 2$) might occur when modelling real-time constraints.

5 Zeno-timelocks: Practice

This section presents a brief description of our zeno-timelock checker, and illustrates its application on a concrete example: the widely used Ethernet protocol CSMA/CD (Carrier Sense Multiple Access with Collision Detection). We will show how a seemingly reasonable timed automata specification of the CSMA/CD (which is inspired by a previous Kronos specification of the same problem [10]) suffers from timelocks. In particular, the example will show how a zeno-timelock occurs which also, and perhaps more dangerously, hides the occurrence of a time-actionlock.

It is important to mention that one can verify that a given UPPAAL specification is free from actionlocks by checking satisfiability of $A[\text{not deadlock}]$ (`deadlock` is a UPPAAL predefined formula). However, this check cannot distinguish between pure-actionlocks and time-actionlocks. Detection of timelocks, and in particular of zeno-timelocks, is hard in UPPAAL. Zeno-timelocks can only be detected with the help of a *test automaton*. Fig. 2a shows a test automaton in UPPAAL; this is included in the original system as a new autonomous component, it does not synchronise with any other component, and τ is a clock local to the automaton. The original system would be free from timelocks if a state where $\tau==1$ can be reached from every state where $\tau==0$, i.e. if the system can always pass time by 1 unit (clocks in UPPAAL can be compared only with integer constants). The formula $(\tau==0) \rightarrow (\tau==1)$ can be written in UPPAAL to verify that the system is free from timelocks. However, this approach provides a sufficient but not necessary condition: if the formula is satisfied the system is guaranteed to be timelock-free, but the formula may be unsatisfiable in some timelock-free systems. The formula $(\tau==0) \rightarrow (\tau==1)$ is actually implementing $A[(\tau==0) \Rightarrow A \langle \tau==1 \rangle]$, which is satisfiable only

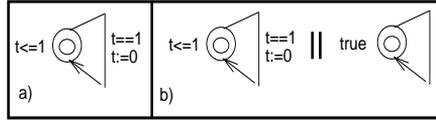


Fig. 2. Test automaton (a) and a timelock-free system (b)

if for every $(t=0)$ -state, a $(t=1)$ -state is reachable in *every* possible run from the $(t=0)$ -state. But this condition is too strong; a system with a zeno run but that is free from timelocks will falsify $A[]((t=0) \Rightarrow A\langle t=1 \rangle)$ as the zeno run is a path where a $(t=1)$ -state is unreachable. Fig. 2b shows such a system. In fact, a system is timelock-free if there exists *at least* one run starting at every $(t=0)$ -state where a $(t=1)$ -state is reachable. It turns out that this condition can be expressed by the formula $A[]((t=0) \Rightarrow E\langle t=1 \rangle)$, but unfortunately such a formula cannot be written in UPPAAL. Also, reachability analysis may suffer from state-explosion, and should a zeno-timelock occur in the system, the trace which witnesses the failure of $(t=0) \rightarrow (t=1)$ may not be meaningful enough to discover the cause of the timelock.

We claim that in many situations our tool will more conveniently assist the user to find zeno-timelocks. Like the test automaton, our tool implements a strategy which is sufficient to detect that a system is free from zeno-timelocks, but does not necessarily imply that a system contains zeno-timelocks. Unlike the test-automaton strategy, however, the analysis here is syntactic (and therefore it may be considerably less demanding than reachability), and it also identifies potential causes of zeno-timelocks directly on the automata structure. This section will then show how our tool complements UPPAAL in the verification of the protocol.

5.1 A zeno-timelock checker

The tool receives a timed automata specification (as an XML file) as input and returns a list of loops which can potentially cause zeno-timelocks. The tool is intended to be integrated with UPPAAL; the user can take advantage of UPPAAL's graphical interface and its rich modelling language to specify the system. This specification is stored by UPPAAL as an XML file which can be input to the zeno-timelock checker. Basically, a cycle-detection algorithm is performed on this specification to discover all structural loops. A second stage determines which loops are strongly non-zeno. Then, loops are matched according to their half-actions; this stage returns a list of matching pairs and a second list of loops which do not contain half-actions. Finally, Lemma 4.3 is applied to both lists to return unsafe pairs/single loops: a pair is unsafe if neither loop is strongly non-zeno, similarly a non-synchronising loop is unsafe

if it is not strongly non-zeno.

5.2 The CSMA/CD protocol

The CSMA/CD (Carrier Sense Multiple Access with Collision Detection) protocol is widely used on Ethernet networks, where the protocol controls the transmission of data between stations sharing a common medium. The following description is mainly taken from [8].

A station wishing to transmit a frame first listens to the medium to determine if another transmission is in progress. If the medium is idle, the station begins to transmit. Otherwise the station continues to listen until the medium is idle, then it begins to transmit immediately. It may happen that two or more stations begin to transmit at about the same time. If this happens, there will be a collision and the data from both transmissions will be garbled and not received successfully. If such a collision is detected during transmission, the station transmits a brief jamming signal (to ensure that all stations know that there has been a collision) and then it ceases transmission. After transmitting the jamming signal, the station waits a random amount of time and then attempts to retransmit the frame.

Collisions can only occur when more than one station begins transmitting within a short time (the period of the propagation delay). If a station attempts to transmit a frame and there are no collisions during the time it takes for the leading edge of the packet to propagate to the farthest station, then there will be no collision for this frame because all other stations are now aware of the transmission (i.e. the medium will be found busy). Secondly, the time needed to detect a collision is no greater than twice the propagation delay.

Fig. 3(a-c) show part of a possible CSMA/CD specification in UPPAAL. Only two stations have been considered in this specification, **Station1** (Fig. 3a) and **Station2** (similar to Fig. 3a modulo renaming). The main role of **Station1** is to model the transmission of frames and retransmission in case collision has been detected. **Medium** (Fig. 3c) will model the state of the medium, i.e. whether collisions have been detected and the broadcast of the jamming signal should any collision occur. Both **Station1** and **Medium** have temporal constraints derived from either the end-to-end propagation delay ($26 \mu\text{s.}$) or the frame-transmission time ($782 \mu\text{s.}$)⁸. We have also included the automaton **UpperLayer1** (Fig. 3b) to model a client layer which uses the protocol service in the station (**UpperLayer2** is similar). It simply provides frames to the protocol layer, acknowledges ongoing transmission and successful termination.

Automaton **Station1** starts in state **Idle**, waiting for **UpperLayer1** to

⁸ Constants respect the IEEE 802.3 standard (Ethernet CSMA/CD)

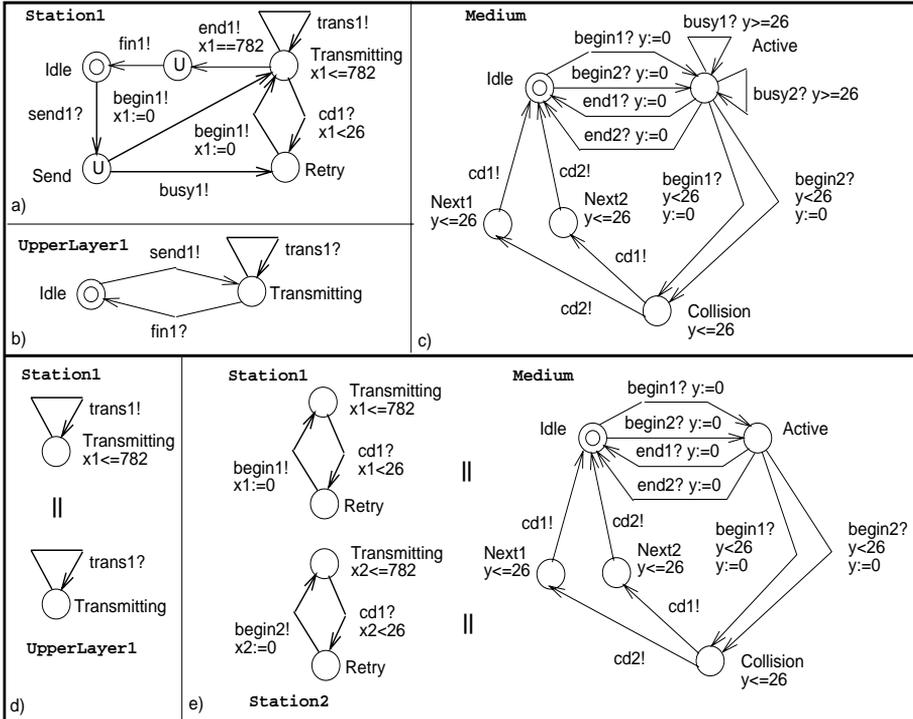


Fig. 3. CSMA/CD in UPPAAL (a,b,c) and unsafe loops (d,e)

send a new frame (**send1?**). If this happens **Station1** moves to **Send**, which is an urgent state: the station may find that either the medium is idle, and so the transmission of the new frame can start immediately (**begin1!**), or that the medium is busy and so the station has to wait (**busy1!**). Urgent states must be left immediately after they are entered; immediate interleaving of actions is permitted but outgoing transitions will be taken with no delay. State **Transmitting** denotes that a transmission has started. Transmission of a complete frame takes 782 μ s, which is modelled both by the invariant $x1 \leq 782$ and the guard $x1 == 782$ on transition **end1!**. Immediately after ending a transmission, a signal **fin1!** is sent to the upper layer to indicate that transmission is completed. While transmission is taking place a signal **trans1!** might be sent to the upper layer to indicate this fact. A collision with another station may occur in **Transmitting**, in which case the jamming signal **cd1?** will be detected. The related guard $x1 < 26$ denotes that no collision can occur after 26 μ s. have passed since a station begun sending a frame. State **Retry** denotes that a collision indeed occurred and that the station is waiting to attempt a retransmission (**begin1!**). The station will remain in **Retry** if a retransmission attempt finds a busy medium; transition **begin1?** is not enabled in such a situation ($x1 \geq 26$).

The **Medium** starts in **Idle**, waiting for stations to begin their transmissions (**begin1?/begin2?**); then it moves to **Active** and clock y is reset. State **Active** denotes that a station is currently using the medium. In **Active**, y denotes the time elapsed since the station begun its transmission. Transitions **busy1?/busy2?** denote that stations can already acknowledge that the medium is in use and thus, that no new transmission is yet possible. The guard $y \geq 26$ in these transitions denote that, in the worst case, a second station cannot acknowledge that the medium is busy before 26 μ s. (the propagation delay) have passed since the first station begun its transmission. State **Collision** denotes that a collision has happened, and that the jamming signal is about to reach the stations. The **Medium** moves from **Active** to **Collision** through **begin1?/begin2?** happening at $y < 26$, i.e. a second station has started transmitting a frame before it could acknowledge that the medium was already in use. In **Collision**, y denotes the time elapsed since a collision occurred; notice that y is reset when the second transmission begins while **Medium** is in **Active** (to simplify matters, we have assumed that a collision occurs as soon as this second transmission begins). The group of transitions **cd1!-Next2-cd2!** and **cd2!-Next1-cd1!** model the jamming signal reaching **Station1** and **Station2**, in any order. Moreover, the invariants $y < 26$ in **Collision**, **Next1/Next2** indicate that the jamming signal will reach the stations not later than 26 μ s. after the collision.

5.3 Verification

We will see how the inclusion of automaton **UpperLayer1** (**UpperLayer2**) disguises a time-actionlock which is already present in the protocol specification, making it undetectable to UPPAAL. In fact, this hidden time-actionlock results in a zeno-timelock which our tool will help to identify.

We begun our verification by checking actionlock-freedom; UPPAAL finds that **A[]not deadlock** is satisfiable in our CSMA/CD specification. We then use our zeno-checker which discovers that a number of synchronising pairs of loops are unsafe and could thus potentially cause zeno-timelocks. These unsafe loops correspond to the interaction between **Station1** and **UpperLayer1** (Fig. 3d), respectively between **Station2** and **UpperLayer2** (not shown), and between **Station1**, **Station2** and **Medium** (Fig. 3e).

Fig. 3e shows a number of loops which could potentially cause zeno-timelocks. We describe these loops below; we use $\langle s1, m, s2 \rangle$ to denote a state in the product automaton where $s1$, m and $s2$ are respectively states in **Station1**, **Medium** and **Station2**. **R**, **I**, **T**, **A**, **C**, **N1** and **N2** respectively denote states **Retry**, **Idle**, **Transmitting**, **Active**, **Collision**, **Next1** and **Next2**. Complete actions **begin1**, **begin2**, **cd1** and **cd2** result from synchronisation

between the corresponding half-actions.

- (i) $\langle R, I, R \rangle$, `begin1`, $\langle T, A, R \rangle$, `begin2`, $\langle T, C, T \rangle$, `cd1`, $\langle R, N2, T \rangle$, `cd2`, $\langle R, I, R \rangle$
- (ii) $\langle R, I, R \rangle$, `begin2`, $\langle R, A, T \rangle$, `begin1`, $\langle T, C, T \rangle$, `cd1`, $\langle R, N2, T \rangle$, `cd2`, $\langle R, I, R \rangle$
- (iii) $\langle R, I, R \rangle$, `begin1`, $\langle T, A, R \rangle$, `begin2`, $\langle T, C, T \rangle$, `cd2`, $\langle T, N1, R \rangle$, `cd1`, $\langle R, I, R \rangle$
- (iv) $\langle R, I, R \rangle$, `begin2`, $\langle R, A, T \rangle$, `begin1`, $\langle T, C, T \rangle$, `cd2`, $\langle T, N1, R \rangle$, `cd2`, $\langle R, I, R \rangle$

These loops correspond to situations in which stations continue to retransmit their frames too soon, therefore colliding again after every attempt. They are considered unsafe because there are no structural conditions ensuring that time will pass in every iteration; i.e. they are not strongly non-zeno (notice in Fig. 3e that clocks are reset but there are no guards with non-zero lower-bounds). In other words, these loops allow zeno runs corresponding to retransmissions following collisions with no delay. However the composite state $\langle R, I, R \rangle$, whose invariant is `True` (because invariants in `Retry` and `Idle` are `True`), is included in every loop. Therefore every loop satisfies the *location non-urgency* property presented in Section 4, and thus they do not cause zeno-timelocks (see Lemma 4.6). Intuitively, there will always exist some infinite execution of every loop which can pass time in state $\langle R, I, R \rangle$.

Now we will focus our attention on the unsafe loop in `Station1` (Fig. 3d); a zeno-timelock would occur in state `Transmitting` (Fig. 3a) if `trans1!` is the only enabled transition at `x1==782`. If this is the case then the invariant in `Transmitting` will make this transition urgent, and so it will be infinitely taken without time passing at all. Should such a zeno-timelock occur in this specification, an actionlock should occur in a specification where transition `trans1!` is removed. As a rationale for this conclusion one has to consider that for a zeno-timelock to involve `trans1!`, this has to be the only transition enabled by `Station1` at `x1==782`. Therefore UPPAAL should be able to detect a “hidden” actionlock if `trans1!` is removed from the specification. This does indeed turn out to be the case, specifically if `trans1!` is removed, UPPAAL detects an actionlock in the resulting system. This is caused by an error in the guard of transition `cd1?` in `Station1`, `x1<26` (note: [10] highlighted the same error). This guard expresses the fact that if there is a collision, this cannot occur after 26 μ s. have passed since `Station1` started transmitting a frame. But 26 μ s. happens to be too small an upper bound for collision detection, as the following scenario illustrates. This scenario is set with `Station1` starting the transmission, a similar scenario can be described for `Station2`.

- (i) `Station1` starts transmitting a frame, therefore `Station1` moves to state `Transmitting` and `Medium` moves to `Active`.

- (ii) **Station2** starts transmitting a frame just before $26 \mu\text{s}$ have passed since **Station1** started transmitting. This means that **Station2**, because of the propagation delay, has not yet been able to detect that the medium is in use. In terms of the protocol specification, notice that in **Medium**, transition **begin2!** can be taken in state **Active** as long as $y < 26$. At this point, **Station1** remains in **Transmitting**, **Station2** has changed to **Transmitting** and **Medium** has changed to **Collision**. Also, it is important to see that the value of clock **x1** is just about to become 26, and that both **x2** and **y** have been reset.
- (iii) Based on the previous observations, and given the invariant $y < 26$ in state **Collision**, notice that it is possible for **x1** to progress to $26 < x1 < 52$. But then the transition **cd1!** in **Collision** will not be able to synchronise with **cd1?** in **Station1**, as the latter is constrained to $x1 < 26$. Should this happen, transition **cd2!** can still be taken to **Next1**, but here again **cd1!** cannot be taken. It is evident, then, that no action will be enabled in the system while **Medium** remains in **Next1**. Furthermore, the invariant $y < 26$ in **Next1** will also prevent time from diverging, raising a time-actionlock when the value of **y** reaches 26.

This time-actionlock shows that the guard $x1 < 26$ in transition **cd1?** in **Station1** (and respectively in **Station2**) should be modified to account for a bigger delay, i.e. it should be $x1 < 52$. This is saying that after a transmission has started the jamming signal could be detected up to $52 \mu\text{s}$. later, that is, twice the propagation delay (see [8] for a detailed explanation). Also, notice that the timelock in this specification resulted in a zeno-timelock in the original specification (i.e. before **trans1!** was removed). When **Medium** is in state **Collision** and $y = 26$, and **Station1** and **Station2** are in state **Transmitting**, transition **trans1!** (**trans2!**) will be infinitely taken while time is prevented from passing (since synchronisation is always possible with **UpperLayer1/UpperLayer2**).

Now, if we correct the specification with the proper delay ($x1 < 52$ in **cd1?**), we can verify that it is free from actionlocks (and thus from time-actionlocks) using UPPAAL's **A[]not deadlock** formula. Since now the time-actionlock in question no longer arises, the loop **trans1!** in the original specification (Fig. 3a) will not cause a zeno-timelock. Time will not be prevented from passing in **Next1/Next2** (Fig. 3c), so the system is allowed to evolve normally and after a collision the stations will move from **Transmitting** to **Retry**, i.e. **trans1!** in **Transmitting** will no longer be enabled.

To clarify then we have taken a specification of the CSMA/CD and attempted to show it is free from zeno-timelocks. We have applied the only

check available in UPPAAL that can throw light on zeno-timelock freedom: checking `A[]not deadlock`. This formula was found to be true, i.e. the system was safe (from an UPPAAL perspective). However we then applied our zeno-timelock checker, which identified a number of potentially unsafe loops (in the sense that they could possibly yield zeno-timelocks). Furthermore, one of these was indeed found to cause a zeno-timelock (note, since an action is always offered, `A[]not deadlock` cannot detect such a zeno-timelock). We thus removed the offending loop from the system and found that the zeno-timelock was indeed “covering” a time-actionlock, which could of course, be detected by UPPAAL once the zeno-loop was removed. The system was corrected to remove this time-actionlock and by so doing we justified that the original offending loop was no longer causing a zeno-timelock.

6 Conclusions

We have identified different types of timelocks which may arise in Timed Automata, and provided formal definitions for each one of them. One of the main contributions of this paper is a new procedure to check whether a system is free from zeno-timelocks. We have refined the syntactic check based on strong non-zenoness, suggested in [9], by carefully analysing the relationship between strong non-zenoness and synchronisation. This allows for the recognition of a wider class of safe (i.e. zeno-timelock free) systems. We have also presented a tool which we have developed which implements this check on timed automata, and in particular UPPAAL specifications. Moreover, this tool can be used to complement the verification capabilities offered by UPPAAL.

We have illustrated the use of our tool on a real-life case-study, the CSMA/CD protocol. We have specified this protocol in UPPAAL and introduced both communication with an upper layer and an incorrect bound for one of the automaton transitions. This was intended to show how hard the detection of some modelling errors can be. The flaw in our specification resulted in a zeno-timelock, which UPPAAL cannot properly detect. The detection of timelocks with the help of a test automaton depends upon a reachability formula which expresses sufficient but not necessary conditions, and reachability analysis may be computationally expensive. Our tool also helped to identify the zeno-timelock in our case-study, showing which structural loops were potentially unsafe (extending ideas devised by Tripakis). The tool is also based upon sufficient but not necessary conditions, however the analysis is syntactic and therefore less demanding than reachability, and it can directly point to sources of zeno-timelocks. Therefore, even if some zeno-timelock free systems are not guaranteed to be so by the sufficient-only conditions, the method pre-

sented in this paper is still useful for narrowing the analysis to specific parts of the model. We are currently trying to develop sufficient and necessary conditions for the detection of zeno-timelocks at the level of the product automaton. Also, we are considering new ways of exploiting the relationship between strong non-zenoness and synchronisation which may further extend the verification scope of our sufficient-only method.

References

- [1] Bornot, S. and J. Sifakis, *On the composition of hybrid systems*, in: *Hybrid Systems: Computation and Control*, LNCS 1386, 1998, pp. 49–63.
- [2] Bowman, H., *Modelling timeouts without timelocks*, in: *ARTS'99, Formal Methods for Real-Time and Probabilistic Systems, 5th International AMAST Workshop*, LNCS 1601 (1999), pp. 335–353.
- [3] Bowman, H., *Time and action lock freedom properties for timed automata*, in: S. K. M. Kim, B. Chin and D. Lee, editors, *FORTE 2001, Formal Techniques for Networked and Distributed Systems* (2001), pp. 119–134.
- [4] Bowman, H., G. Faconti, J.-P. Katoen, D. Latella and M. Massink, *Automatic verification of a lip synchronisation algorithm using UPPAAL*, *Formal Aspects of Computing* **10** (1998), pp. 550–575.
- [5] Bowman, H., R. Gomez and L. Su, *How to stop time stopping (preliminary version)*, Technical Report 9-04, Computing Laboratory, University of Kent, UK (2004).
- [6] Milner, R., “Communication and Concurrency,” Prentice-Hall, 1989.
- [7] Pettersson, P. and K. G. Larsen, *UPPAAL2K: Small Tutorial*, *Bulletin of the European Association for Theoretical Computer Science* **70** (2000), pp. 40–44.
- [8] Stallings, W., “Data & Computer Communications,” Prentice Hall, 2000, 6th. edition.
- [9] Tripakis, S., *Verifying progress in timed systems*, in: *ARTS'99, Formal Methods for Real-Time and Probabilistic Systems, 5th International AMAST Workshop*, LNCS 1601 (1999).
- [10] Yovine, S., *Kronos: A verification tool for real-time systems*, *Springer International Journal of Software Tools for Technology Transfer* (1997).

A Proofs

The following notation (where $A \in TA$ refers to a single automaton, and $A[1], \dots, A[n] \in TA$ refers to a network of automata) will be used throughout this section⁹:

- $Loops(A)$ is the set of all structural loops in A .
- An *edge* $e = (a, g, r)$ is the edge-labelling of $l_1 \xrightarrow{a,g,r} l_2 \in T$.

⁹ Elements which have been previously introduced in the paper are recalled here just for the sake of completeness.

- $Edges(A)$, $Edges(lp)$ denote, respectively, the set of edges in A and in loop lp .
- $Loc(lp)$ is the set of locations in loop lp .
- A *half loop* is a loop which contains at least one edge labelled with a half action, i.e. $lp \in Loops(A)$ s.t. $\exists (a, g, r) \in Edges(lp)$. $a \in HA$. A *complete loop* is a loop which is not a half loop.
- Two *synchronising loops* lp_1, lp_2 , denoted $sync(lp_1, lp_2)$, are half loops in different component automata with matching half actions, i.e. $\exists A[i], A[j]$ ($i \neq j$), $lp_1 \in Loops(A_i), lp_2 \in Loops(A_j), e_1 \in Edges(lp_1), e_2 \in Edges(lp_2)$. $e_1 = (a?, g_1, r_1) \wedge e_2 = (a!, g_2, r_2)$.
- A *composite edge* is any $e = (a, g_1 \wedge g_2, r_1 \cup r_2) = e_1 || e_2$, where $e_1 = (a?, g_1, r_1)$ and $e_2 = (a!, g_2, r_2)$.
- A *composite loop*, denoted $comp(lp)$, is s.t. $\exists A[i], A[j], e_i \in Edges(A[i]), e_j \in Edges(A[j])$. $e_i || e_j \in Edges(lp)$.
- Given two loops lp_1, lp_2 we say that lp_1 is *included* in lp_2 , denoted $lp_1 \subseteq lp_2$, if $Loc(lp_1) \subseteq Loc(lp_2)$ and $\forall e \in Edges(lp_1). (e \in Edges(lp_2) \vee \exists A[i], e_i \in Edges(A[i]). e || e_i \in Edges(lp_2))$
- HL is the set of all pairs of synchronising loops in $A[1], \dots, A[n]$, i.e. $HL = \{(lp_i, lp_j) \mid \exists A[i], A[j], (i \neq j). lp_i \in Loops(A[i]) \wedge lp_j \in Loops(A[j]) \wedge sync(lp_i, lp_j)\}$.
- CL is the set of all complete loops in $A[1], \dots, A[n]$, i.e. $CL = \{lp \mid \exists A[i]. lp \in Loops(A[i]) \wedge \forall (a, g, r) \in Edges(lp). a \in CA\}$.

Lemma A.1 *Every composite loop contains at least two synchronising loops. Formally, $\forall lp \in Loops(|A)$. $comp(lp) \Rightarrow \exists lp_i \in Loops(A[i]), lp_j \in Loops(A[j]). sync(lp_i, lp_j) \wedge lp_i \subseteq lp \wedge lp_j \subseteq lp$*

Proof. Automata theory gives us the necessary clues,

- 1 Given l_1, l_2 as two locations in $|A$, any path $l_1 \xrightarrow{*} l_2$ in $|A$ must include a path $l_1[i] \xrightarrow{*} l_2[i]$ from every $A[i]$.
- 2 Moreover, if there is a composite edge $e = e_i || e_j$ in $l \xrightarrow{*} l'$, e_i must be an edge in $l_1[i] \xrightarrow{*} l_2[i]$ and e_j an edge in $l_1[j] \xrightarrow{*} l_2[j]$.

Therefore, since every loop is by definition a path, the lemma easily follows. \square

Lemma A.2 *Clock reset and “lower-boundedness” are preserved in composite edges. Formally, let $e_1 = (a_1, g_1, r_1)$, $e_2 = (a_2, g_2, r_2)$, $e_3 = (a_3, g_3, r_3)$ denote three edges, $c \in C$ a given clock and $\epsilon \in \mathbb{R}^+$ a given constant. If $e_3 = e_1 || e_2$ then*

$$1 \quad c \in r_1 \cup r_2 \Rightarrow c \in r_3$$

$$2 \ (g_1 \Rightarrow c > \epsilon) \vee (g_2 \Rightarrow c > \epsilon) \Rightarrow (g_3 \Rightarrow c > \epsilon)$$

Proof. Follows from the definition of edge composition. \square

Lemma A.3 *Loop inclusion preserves strong non-zenoness. Formally, let lp_1, lp_2 be two loops s.t. $lp_1 \subseteq lp_2$. If lp_1 is strongly non-zero then lp_2 is also strongly non-zero.*

Proof.

- 1 $lp_1 \subseteq lp_2$ (hyp.).
- 2 lp_1 is strongly non-zero (hyp.).
- 3 $\exists e_i, e_j \in Edges(lp_1), c \in \mathbb{C}, \epsilon \in \mathbb{R}^+, a, a' \in \mathbb{A}. e_i = (a, g_i, r_i) \wedge e_j = (b, g_j, r_j) \wedge c \in r_i \wedge g_j \Rightarrow c > \epsilon$ (by 2 and def. of strong non-zenoness).
- 4 $e_i \in Edges(lp_2) \vee \exists A[k], e_k \in Edges(A[k]), e \in Edges(lp_2). e = e_k || e_i$ (by 1, 3 and def. of edge composition).
- 5 $e_j \in Edges(lp_2) \vee \exists A[k], e_k \in Edges(A[k]), e \in Edges(lp_2). e = e_k || e_j$ (by 1, 3 and def. of edge composition).
- 6 $\exists e_i, e_j \in Edges(lp_2), c \in \mathbb{C}, \epsilon \in \mathbb{R}^+, a, a' \in \mathbb{A}. e_i = (a, g_i, r_i) \wedge e_j = (a', g_j, r_j) \wedge c \in r_i \wedge g_j \Rightarrow c > \epsilon$ (by 3, 4 and Lemma A.2)
- 7 lp_2 is strongly non-zero (by 6 and def. of strong non-zenoness). \square

Lemma A.4 *If at least one loop in every element of HL is strongly non-zero, then all composite loops in $|A$ are strongly non-zero.*

Proof.

- 1 $lp \in Loops(|A), comp(lp)$ (hyp.).
- 2 $\exists A[i] \neq A[j], lp_i \in Loops(A[i]), lp_j \in Loops(A[j]). sync(lp_i, lp_j) \wedge lp_i \subseteq lp \wedge lp_j \subseteq lp$ (by Lemma A.1).
- 3 $(lp_i, lp_j) \in HL$ (by 2 and def. of HL).
- 4 Suppose lp_i strongly non-zero (hyp.).
- 5 lp is strongly non-zero (by 2 and Lemma A.3). \square

Lemma 4.3 *If (at least) one loop in every pair of HL is strongly non-zero and all loops in CL are strongly non-zero then the product automaton $|A$ is also strongly non-zero and thus free from zeno-timelocks.*

Proof.

- 1 Consider any loop $lp \in Loops(|A)$.

- 2 lp is s.t. either $comp(lp)$ or there exists a complete loop $lp_i \in CL$ s.t. $lp_i \subseteq lp$ (by automata theory).
- 3 We know that at least one loop in every element of HL is strongly non-zero, and that all loops in CL are strongly non-zero (hyp.).
- 4 If $comp(lp)$ then lp is strongly non-zero (by hyp. and Lemma A.4).
- 5 If $lp_i \subseteq lp$, $lp_i \in CL$ then lp is strongly non-zero since lp_i is strongly non-zero (by hyp. and Lemma A.3).
- 6 Therefore all loops in $|A$ are strongly non-zero, and so $|A$ is free from zeno-timelocks (by Lemma 4.2).

□

Lemma 4.6 *If $A \in TA$ is either location non-urgent or reset non-urgent, then it is also free from zeno-timelocks.*

Proof.

- 1 Consider $A \in TA$ as either location non-urgent or reset non-urgent (hyp.).
- 2 Consider a given state s , and ρ an infinite run starting at s (hyp.).
- 3 ρ must visit a given loop lp in A an infinite number of times (by def. of infinite run).
- 4 If A is location non-urgent there must be a location l in the loop lp where either $I(l) : True$ or all clocks in the invariant are unbounded (by def. of location non-urgency).
- 5 Any execution of A which reaches l can wait indefinitely in l (by 4). Note that because we are assuming strong invariants, invariant expressions such as $I(l) : x > c$ (where $x \in Clocks(I(l))$ and $c \in \mathbb{R}^{+0}$) do not force any execution to leave location l immediately. Moreover, l is only reachable if at least c time units have passed since the last time x was reset.
- 6 The location l is reachable in ρ , i.e. $\rho = s \xrightarrow{*} [l, v] \xrightarrow{*}$, where $v \in \mathbb{V}_C$ (by 3).
- 7 There exists a time-unbounded run ρ' starting at s , i.e. $\rho' = s \xrightarrow{*} [l, v] \xrightarrow{\infty}$. Therefore, if A is location non-urgent then s cannot be a zeno-timelock (by 5, 6 and def. of zeno-timelock).
- 8 If A is reset non-urgent there must be a location l in the loop lp where at least a clock in the invariant has a non-zero lower-bound, say $d \in \mathbb{R}^+$, and it is reset in a given transition i of lp (by def. of reset non-urgency).
- 9 Any execution of A which takes transition i and then reaches location l must have elapsed at least d time units between these two events (by 8

and assuming strong invariants).

- 10 ρ takes transition i and visits location l an infinite number of times (by 3 and 8).
- 11 ρ accumulates an infinite number of d time units, so $delay(\rho) = \infty$. Therefore, if A is reset non-urgent then s cannot be a zeno-timelock (by 9, 10 and def. of zeno-timelock).
- 12 A is free from zeno-timelocks (by 7 and 11).

□