

Realizability of Real-Time Logics

L. Doyen¹, G. Geeraerts¹, J.-F. Raskin¹, and J. Reichert²

¹ Département d'Informatique, Université Libre de Bruxelles (U.L.B.)

² École Normale Supérieure de Cachan

Abstract. We study the realizability problem for specifications of reactive systems expressed in real-time linear temporal logics. The logics we consider are subsets of MITL (Metric Interval Temporal Logic), a logic for which the satisfiability and validity problems are decidable, a necessary condition for the realizability problem to be decidable. On the positive side, we show that the realizability of LTL extended with past real-time formulas is decidable in $2EXPTIME$, with a matching lower bound. On the negative side, we show that a simple extension of this decidable fragment with future real-time formulas leads to undecidability. In particular, our results imply that the realizability problem is undecidable for ECL (Event Clock Logic), and therefore also for MITL.

1 Introduction

The *satisfiability* and *model-checking problems* for real-time temporal logics have been widely studied since the nineties [4, 12]. The main application of these problems is the verification of reactive systems: given a model of the system and of its environment, one can check whether the parallel composition of the two models satisfies a specification given by a real-time logic formula. This well-established procedure applies to *closed models* obtained when both the system and the environment are fully specified.

However, in the design of real-time reactive systems, such a precise model of the reactive system is usually difficult to construct manually; and on the other hand, the environment may be only partially known, especially in the early stages of development. Therefore, it is natural to consider the problem of the automatic synthesis of a behavior policy for the reactive system that would be correct by construction with respect to the specification. This problem is usually formalized as a two-players game, in which Player 1 controls the execution of the system, and Player 2 controls the execution of environment. The specification is encoded as the winning condition for Player 1 in the game. Roughly speaking, the behaviors of Player 1 represent all possible models for the system, and computing a winning strategy for Player 1 amounts to selecting one model which is guaranteed to be correct whatever the environment does.

In the setting of timed systems, most of the previous works³ have considered games played on *deterministic timed automata*, whose set of edges is partitioned into those controlled by Player 1 and those controlled by Player 2. The winning condition is a simple safety objective (some set of locations should be avoided, no matter what the environment does), or more generally, an ω -regular objective defined as a parity condition over the locations of the automaton.

In this paper, we consider an abstract definition of two-players timed games with a winning condition expressed by a real-time temporal logic formula. Consider a finite

³ With the notable exception of [17] see the ‘related works’ paragraph.

set Σ_1 of actions controlled by Player 1, and a finite set Σ_2 of actions controlled by Player 2. Let φ be a real-time temporal logic formula defining a set of timed words over $\Sigma = \Sigma_1 \cup \Sigma_2$. The timed game is played for infinitely many rounds as follows. In each round, the players announce (simultaneously and independently of each other) a pair (Δ, α) consisting of a delay $\Delta \in \mathbb{R}^{\geq 0}$ and an action α from their set of controllable actions. The player who has announced the shortest delay is allowed to play its action after the corresponding delay, and then the next round starts. The outcome of the game is an infinite timed word. Player 1 wins the game if the outcome satisfies the formula φ . Note that no game graph needs to be provided in this abstract definition. The problem to decide whether Player 1 has a strategy to win the game regardless of the choices of Player 2 is called the *realizability problem*, borrowing the terminology introduced for LTL (Linear Temporal Logic) [23]. In a variant of this problem [9], one asks that Player 1 wins without announcing a converging sequence of delays (thus without blocking time), i.e., the outcome has to be either time-diverging and then belong to φ , or time-converging and then Player 1 has announced the shortest delay only finitely often. All results in this paper hold for both variants of the realizability problem.

As it is easy to show that the realizability problem for a logic is at least as hard as both the satisfiability problem and the validity problem for that logic, we need to consider specifications that are expressible in real-time logics for which these two problems are decidable. One of the most natural way to obtain a real-time logic is to equip the modalities of LTL [22] with real-time constraints. For instance, $\diamond_{[a,b]}\varphi$ holds in some position p iff there is a future position p' in which φ holds, and the time elapsed between p and p' is between a and b time units. This extension of LTL is the Metric Temporal Logic (MTL) introduced by Koymans [13]. Unfortunately, it has been shown that the satisfiability problem is undecidable for MTL [5] when interpreted over infinite timed words. However, when prohibiting singular time intervals of the form $[a, a]$, this logic becomes decidable (and is then called Metric Interval Temporal Logic, or MITL) [2]. Another way of obtaining a decidable real-time logic is to extend LTL with new real-time operators, as in the Event Clock Logic (ECL) [25, 24, 12]. In ECL, the operators \triangleright_1 and \triangleleft_1 are introduced, allowing to speak about the *next* (resp. *last*) time a formula will be (was) true. For instance, $\triangleright_{[a,b]}\varphi$ holds in a position p if there exists a future position p' where φ holds, the time elapsed between p and p' is in $[a, b]$, and φ has been false in all positions between p and p' . This is to be contrasted with the intuitive meaning of the MTL formula $\diamond_{[a,b]}\varphi$ which does not constrain the truth value of φ in the interval $[0, a)$. It is known that the expressivity of ECL is subsumed by that of MITL and therefore the satisfiability problem for ECL is decidable [25, 24]. Thus, both MITL and ECL are good candidates for the realizability problem. It is a long-standing open question whether realizability is decidable for MITL. Surprisingly however, a consequence of our results is that the realizability problem for both ECL and MITL is undecidable.

Contributions This paper provides two main theoretical results about the realizability problem for ECL. First, we show that the realizability problem for ECL is undecidable. This result is surprising as this logic can be translated to recursive event-clock automata [25, 24, 3], a determinizable class of timed automata. Unfortunately, those automata are only deterministic in a weak sense, as already noted in [16]: while every infinite word has indeed a unique run in an event-clock automaton, it may be that two timed words with a common prefix (say up to position i) have runs with different prefixes (up to position i). This is due to the fact that runs in event-clock automata constrain

their future by using prophecy clocks. While weak determinism is sufficient to ensure closure under complement for example, our undecidability result formally shows that this notion of determinism is not sufficient to obtain an algorithm for the realizability problem. As ECL is a subset of MITL, this result immediately entails the undecidability of the realizability problem for MITL. Second, we show that LTL extended with the past fragments of ECL (called henceforth LTL_{\triangleleft}), has a decidable realizability problem. We provide a translation of this real-time extension of LTL to classical Alur-Dill deterministic timed automata [1]. Using this translation, we obtain a 2EXPTIME algorithm for the realizability problem, and a matching lower bound since the problem is already 2EXPTIME-hard for LTL.

Related Works As already mentioned, there have been several previous works about timed games, see for instance [19, 7]. In those works, the objectives are specified by deterministic timed automata. We focus here on related works where real-time logics have been used to define the objective of the timed game. In [17], a decidability result is obtained for the realizability problem of bounded-response properties which are expressible a fragment of MTL with future operators. The result holds under the a bounded-variability semantics, i.e., the number of events per time unit is bounded by a constant. In our case, we do not need this hypothesis. Note that under bounded-variability semantics, the full MITL can be translated to deterministic timed automata. In [16], the past fragment of MITL is translated into deterministic timed automata. The logics there are interpreted over finite signals for the purpose of monitoring while our logics are interpreted over infinite timed words for the purpose of realizability. The past fragment of MITL is incomparable with the logic LTL_{\triangleleft} for which we have the decidability result. Note that over finite words, the satisfiability problem for MTL is decidable [20]. Unfortunately, the synthesis problem is in general undecidable even on finite words, but becomes decidable when the resources of the controller are bounded [6].

Remark Due to lack of space, the proofs are given in the appendix.

2 Preliminaries

An *interval* is a nonempty convex subset of the set $\mathbb{R}^{\geq 0}$ of nonnegative real numbers. Intervals may be left-open or left-closed; right-open or right-closed; bounded or unbounded. An interval has one of the following forms: $[a, b]$, $[a, b)$, $[a, \infty)$, $(a, b]$, (a, b) , (a, ∞) , with endpoints $a, b \in \mathbb{N}$ and $a \leq b$. A *word* over a finite alphabet Σ is a (finite or infinite) sequence $w = w_0w_1\dots$ of symbols $w_i \in \Sigma$. We denote by $|w|$ the *length* of w , i.e., the number of symbols in w . A *timed word* over Σ is a pair $\theta = (w, \tau)$ where w is a word over Σ , and $\tau = \tau_0\tau_1\dots$ is a sequence of length $|w|$ of time values $\tau_i \in \mathbb{R}^{\geq 0}$ such that $\tau_i \leq \tau_{i+1}$ for all $0 \leq i < |w|$. We often denote a timed word (w, τ) as a sequence $(w_0, \tau_0)(w_1, \tau_1)\dots$ of symbols paired with their time stamp. An infinite timed word $\theta = (w, \tau)$ is *diverging* if for all $t \in \mathbb{R}^{\geq 0}$, there exists a position $i \in \mathbb{N}$ such that $\tau_i \geq t$.

Automata formalisms. We first define automata on (untimed) words. A (nondeterministic) *finite automaton* over a finite alphabet Σ is a tuple $A = (Q, q_{in}, E, \alpha)$ where Q is a finite set of states, $q_{in} \in Q$ is the initial state, $E \subseteq Q \times \Sigma \times Q$ is a set of transitions, and α is an acceptance condition on transitions. We consider two kinds of acceptance

conditions: the *generalized Büchi condition* when $\alpha \subseteq 2^E$ is a set of sets of transitions, and the *parity condition* with d priorities when $\alpha : E \rightarrow \{0, 1, \dots, d\}$.⁴ The automaton A is *deterministic* if for all states q and all symbols $\sigma \in \Sigma$, there exists $(q, \sigma, q') \in E$ for exactly one $q' \in Q$.

A *run* of a finite automaton A over a word w is a sequence $q_0 w_0 q_1 w_1 q_2 \dots$ such that $q_0 = q_{in}$ and $(q_i, w_i, q_{i+1}) \in E$ for all $0 \leq i < |w|$. For finite runs r , we denote by $\text{Last}(r)$ the last state in r , and for infinite runs r , we denote by $\text{Inf}(r)$ the set of transitions occurring infinitely often in r . An infinite run r is *accepting* according to the generalized Büchi condition α if for all sets of edges $F \in \alpha$, $\text{Inf}(r) \cap F \neq \emptyset$. An infinite run r is *accepting* according to the parity condition α if $\min\{\alpha(e) \mid e \in \text{Inf}(r)\}$ is even. The *language* defined by a finite automaton A , noted $L(A)$, is the set of infinite words on which A has an accepting run.

We next define timed automata over infinite timed words [1]. Let X be a finite set $\{x_1, x_2, \dots, x_n\}$ of variables called *clocks*. An *atomic clock constraint* is a formula of the form $x \in I$ where I is an interval with integer endpoints (and possibly unbounded). A *guard* is a boolean combination of atomic clock constraint. We denote by $\text{Guards}(X)$ the set of all guards on X . A *valuation* for the clocks in X is a function $v : X \rightarrow \mathbb{R}^{\geq 0}$. We write $v \models g$ whenever the valuation v satisfies the guard g . For $R \subseteq X$, we write $v[R := 0]$ for the valuation that assigns 0 to all clocks $x \in R$, and $v(x)$ to all clocks $x \notin R$. For $t \in \mathbb{R}^{\geq 0}$, we write $v + t$ for the valuation that assigns the value $v(x) + t$ to each clock $x \in X$. A *timed automaton* over alphabet Σ and clocks X is a tuple $A = (Q, q_{in}, E, \alpha)$ where Q is a finite set of states, $q_{in} \in Q$ is the *initial state*, $E \subseteq Q \times \Sigma \times \text{Guards}(X) \times 2^X \times Q$ is a set of *transitions*, and α is an acceptance condition, either a *generalized Büchi condition* if $\alpha \subseteq 2^E$, or a *parity condition* with d priorities if $\alpha : E \rightarrow \{0, 1, \dots, d\}$. The timed automaton A is *deterministic* if for every state q and valuation v , for all $\sigma \in \Sigma$, there exists at most one transition $(q, \sigma, g, R, q') \in E$ such that $v \models g$.

A *timed run* r of a timed automaton A over a timed word (w, τ) is an infinite sequence $(q_0, v_0)(w_0, \tau_0)e_0(q_1, v_1)(w_1, \tau_1)e_1 \dots$ such that (i) $q_0 = q_{in}$, (ii) $v_0(x) = 0$ for all $x \in X$, and (iii) for all positions $i \geq 0$, $e_i = (q_i, w_i, g, R, q_{i+1}) \in E$ is such that $v_i + \tau_i - \tau_{i-1} \models g$ and $v_{i+1} = (v_i + \tau_i - \tau_{i-1})[R := 0]$ (assuming $\tau_{-1} = 0$). The definition of *accepting* timed run is adapted from the untimed case. The *timed language* of a timed automaton A , is the set $L(A)$ of timed words on which A has an accepting timed run.

Real-time logics. We consider the logic ECL (Event Clock Logic) and some of its fragments [24, 25, 12]. ECL is an extension of LTL with two real-time operators: the history operator $\triangleleft_I \varphi$ expressing that φ was true for the last time t time units ago for some $t \in I$, and the prediction operator $\triangleright_I \varphi$ expressing that the next time φ will be true is in t time units for some $t \in I$. Given a finite alphabet Σ , the syntax of ECL is the following:

$$\varphi \in \text{ECL} ::= a \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \mathcal{S} \varphi \mid \varphi \mathcal{U} \varphi \mid \triangleleft_I \varphi \mid \triangleright_I \varphi$$

⁴ Acceptance conditions on transitions can be easily transformed into acceptance conditions over states by doubling the state space of the automaton for the generalized Büchi condition and by taking d copies of the state space for the parity condition.

where $a \in \Sigma$ and I is an interval. The models of an ECL formula are infinite timed words. A timed word $\theta = (w, \tau)$ satisfies a formula $\varphi \in \text{ECL}$ at position $i \in \mathbb{N}$, written $\theta, i \models \varphi$, according to the following rules:

- if $\varphi = a$, then $w_i = a$;
- if $\varphi = \neg\varphi'$, then $\theta, i \not\models \varphi'$;
- if $\varphi = \varphi_1 \vee \varphi_2$, then $\theta, i \models \varphi_1$ or $\theta, i \models \varphi_2$;
- if $\varphi = \varphi_1 \mathcal{S} \varphi_2$, then there exists $0 \leq j < i$ such that $\theta, j \models \varphi_2$ and for all $j < k < i$, $\theta, k \models \varphi_1$;
- if $\varphi = \varphi_1 \mathcal{U} \varphi_2$, then there exists $j > i$ such that $\theta, j \models \varphi_2$ and for all $i < k < j$, $\theta, k \models \varphi_1$;
- if $\varphi = \triangleleft_I \varphi'$, then there exists $0 \leq j < i$ such that $\theta, j \models \varphi'$, $\tau_i - \tau_j \in I$, and for all $j < k < i$, $\theta, k \not\models \varphi'$;
- if $\varphi = \triangleright_I \varphi'$, then there exists $j > i$ such that $\theta, j \models \varphi'$, $\tau_j - \tau_i \in I$, and for all $i < k < j$, $\theta, k \not\models \varphi'$;

When $\theta, 0 \models \varphi$, we simply write $\theta \models \varphi$ and we say that θ satisfies φ . We denote by $\llbracket \varphi \rrbracket$ the set $\{\theta \mid \theta \models \varphi\}$ of models of φ . Finally, we define the following shortcuts: $\text{true} \equiv a \vee \neg a$ with $a \in \Sigma$, $\text{false} \equiv \neg \text{true}$, $\varphi_1 \wedge \varphi_2 \equiv \neg(\neg\varphi_1 \vee \neg\varphi_2)$, $\varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2$, $\diamond\varphi \equiv \text{true} \mathcal{U} \varphi$, $\square\varphi \equiv \varphi \wedge \neg\diamond(\neg\varphi)$, $\bigcirc\varphi \equiv \text{false} \mathcal{U} \varphi$, $\ominus\varphi \equiv \text{false} \mathcal{S} \varphi$, and $\hat{\diamond}\varphi \equiv \text{true} \mathcal{S} \varphi$. We also freely use notations like $\geq x$ to denote the interval $[x, \infty)$, and $< x$ for $[0, x)$, etc. in the \triangleleft and \triangleright operators.

Then, we define two fragments of ECL. **PastECL** is the fragment of ECL where the temporal operators speak about the *past* only. A formula φ of ECL is in **PastECL** if there is no occurrence of $\triangleright_I \varphi_1$ and $\varphi_1 \mathcal{U} \varphi_2$ in the subformulas of φ . $\text{LTL}_{\triangleleft}$ is an extension of LTL [22] with the \triangleleft_I operator from ECL, with the restriction that only formulas of **PastECL** appear under the scope of a \triangleleft_I . A formula ψ of ECL is in $\text{LTL}_{\triangleleft}$ if (i) when $\triangleleft_I \varphi_1$ is a subformula of ψ , then $\varphi_1 \in \text{PastECL}$, and (ii) there is no $\triangleright_I \varphi_1$ in the subformulas of ψ . Formally,

$$\begin{aligned} \varphi \in \text{PastECL} &::= a \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \mathcal{S} \varphi \mid \triangleleft_I \varphi \\ \psi \in \text{LTL}_{\triangleleft} &::= a \mid \neg\psi \mid \psi \vee \psi \mid \psi \mathcal{S} \psi \mid \psi \mathcal{U} \psi \mid \triangleleft_I \varphi \end{aligned}$$

The truth value of a formula φ of **PastECL** at a position i in a timed word θ depends only on the events of θ at positions j , with $0 \leq j \leq i$. On the other hand, a formula of $\text{LTL}_{\triangleleft}$ may speak about the future of a word, but not in an untimed fashion, i.e., only using the (untimed) \mathcal{U} operators.

Example 1.

- $\varphi_1 \equiv \square(c \rightarrow \triangleright_{(2,3)} a)$ is a formula of ECL (but neither of $\text{LTL}_{\triangleleft}$, nor of **PastECL**), saying that every c is followed by an a , between 2 and 3 time units.
- $\varphi_2 \equiv \square(c \rightarrow \triangleleft_{(2,3)}(a \wedge \ominus b))$ is a formula of $\text{LTL}_{\triangleleft}$ (but not of **PastECL**) saying that every c has to be preceded, between 2 and 3 time units before, by an a directly precede by a b .
- $\varphi_3 \equiv a \wedge \triangleleft_{(2,3)} c$ is a **PastECL** formula that holds in all positions where an a occurs preceded, between 2 and 3 time units before, by a c .

Timed games and realizability. A *timed game* is a tuple $G = \langle \Sigma_1, \Sigma_2, W \rangle$ where Σ_j is the finite alphabet of Player j ($j = 1, 2$), $\Sigma_1 \cap \Sigma_2 = \emptyset$, and W is a set of timed words over $\Sigma_1 \cup \Sigma_2$, called the winning condition for Player 1. Timed games are played as in [9] but with a trivial game structure.

A timed game is played for infinitely many rounds as follows. In round i ($i \geq 0$), Player 1 chooses a time delay $\Delta_1^i \in \mathbb{R}^{\geq 0}$ and an action $\alpha_1^i \in \Sigma_1$, while independently and simultaneously, Player 2 chooses a time delay $\Delta_2^i \in \mathbb{R}^{\geq 0}$ and an action $\alpha_2^i \in \Sigma_2$. Then, a *play* in G is a timed word $(w_0, \tau_0)(w_1, \tau_1) \dots$ over $\Sigma_1 \cup \Sigma_2$ such that for all $i \geq 0$ (assuming $\tau_{-1} = 0$), $\tau_i = \tau_{i-1} + \min\{\Delta_1^i, \Delta_2^i\}$ and either $w_i = \alpha_1^i$ and $\Delta_1^i \leq \Delta_2^i$, or $w_i = \alpha_2^i$ and $\Delta_2^i \leq \Delta_1^i$. Note that there can be several plays produced by a given sequence of choices of the players, namely if $\Delta_1^i = \Delta_2^i$ for some i . We say that Player j *plays first* in round i if $\alpha^i = \alpha_j^i$, and we denote by $Blameless_1$ the set of timed words over $\Sigma_1 \cup \Sigma_2$ that contain only finitely many letters from Σ_1 .

A timed word θ is *winning* for Player 1 if $\theta \in W$. Let td be the set of diverging timed words on $\Sigma_1 \cup \Sigma_2$. A timed word θ is *td-winning* for Player 1 if $\theta \in WC_1(W) = (W \cap td) \cup (Blameless_1 \setminus td)$, i.e., Player 1 wins because the word is diverging and belongs to W , or because Player 2 is responsible for the convergence of time.

A strategy for player j is a function π that maps every finite timed word $\theta = (w_0, \tau_0)(w_1, \tau_1) \dots (w_n, \tau_n)$ to a pair (Δ_j, α_j) , where $\Delta_j \in \mathbb{R}^{\geq 0}$ and $\alpha_j \in \Sigma_j$. A play $(w_0, \tau_0)(w_1, \tau_1) \dots$ is *consistent* with π (for player j) if for all $i \geq 0$, either $\pi((w_0, \tau_0)(w_1, \tau_1) \dots (w_i, \tau_i)) = (w_{i+1}, \tau_{i+1} - \tau_i)$, or $w_{i+1} \in \Sigma_{3-j}$ and $\Delta_j \geq \tau_{i+1} - \tau_i$ where $(\Delta_j, \cdot) = \pi((w_0, \tau_0)(w_1, \tau_1) \dots (w_i, \tau_i))$.

We denote by $Outcome_j(G, \pi)$ the set of all plays in G that are consistent with π . A strategy π for player 1 is *winning* (resp. *td-winning*) for player 1 if $Outcome_1(G, \pi)$ contains only winning (resp. td-winning) plays. Finally, given a strategy π_1 for Player 1 and a strategy π_2 for Player 2, let $Outcome(G, \pi_1, \pi_2)$ denote the set of all possible plays in G that are consistent with π_1 and π_2 . Note that $Outcome(G, \pi_1, \pi_2)$ is not necessarily a singleton since there is nondeterminism in the game when the same delay is proposed by the two players.

The *realizability problem* (resp. *td-realizability problem*) for a logic L is to decide, given two finite sets Σ_1, Σ_2 and a formula $\varphi \in L$ over $\Sigma_1 \cup \Sigma_2$, whether Player 1 has a winning (resp. td-winning) strategy in the timed game $\langle \Sigma_1, \Sigma_2, \llbracket \varphi \rrbracket \rangle$.

Example 2. Consider the game $G_e = \langle \Sigma_1, \Sigma_2, \llbracket \varphi_e \rrbracket \rangle$, where $\Sigma_1 = \{a, b\}$, $\Sigma_2 = \{c, d\}$ and $\varphi_e \equiv \varphi_H \rightarrow \varphi_C$ where $\varphi_H \equiv \Box \left(c \rightarrow (\bigcirc(\Box \neg c) \vee (\neg c) \mathcal{U} a) \right)$ and $\varphi_C \equiv \Box \left((c \rightarrow \diamond a) \wedge (a \rightarrow (\neg \diamond c \vee \triangleleft_{(2,3)} c)) \right)$. In this game, Player 2 makes requests by playing c 's, and Player 1 has to acknowledge the request by outputting a 's. The assumption φ_H prevents Player 2 to issue a second request before the first one has been acknowledged. On the other hand, the condition φ_C forces Player 1 to acknowledge every request within 2 to 3 time units. Moreover, Player 1 can play b 's and Player 2 can play d 's freely.

In this game, Player 1 has a td-winning strategy but no winning strategy: every time a c is played at time t_c , Player 1 proposes to play $(a, 2.5 - (t - t_c))$ at every time stamp t until an a has been played. More precisely, for a prefix $\theta = (w, \tau)$ of length ℓ : $\pi(\theta) = (a, 2.5 - (\tau(\ell - 1) - \tau(j)))$ if there exists $i < \ell - 1$ such that $w(i) = c$ and $w(j) \neq a$ for all $i < j \leq \ell - 1$. Otherwise, $\pi(\theta) = (b, 1)$. Thus, either Player 1 eventually plays first and an a is played 2.5 time units after the c . Or Player 1 never

plays first again, and Player 2 is blocking the time. In both case, this is a td-winning play for Player 1. However, this is not a winning play, and there cannot be any winning play for Player 1, since she cannot prevent Player 2 from blocking the time after a c has been played.

Lossy 3-counter machines A deterministic lossy 3-counter machine (3CM) [15] is a tuple $M = \langle c_1, c_2, c_3, Q, q_{in}, \delta \rangle$ where c_1, c_2 , and c_3 are three nonnegative counters, Q is a finite set of states, $q_{in} \in Q$ is the initial state, and $\delta : Q \mapsto I$ is the transition function where I is a finite set of instructions of the form c_i++ ; goto q or if $c_i \neq 0$ then c_i-- ; goto q else goto q' or halt, for $i \in \{1, 2, 3\}$ and $q, q' \in Q$.

A configuration of a 3CM M is a tuple $\gamma = (q, \nu^1, \nu^2, \nu^3)$ where $q \in Q$ and $\nu^1, \nu^2, \nu^3 \in \mathbb{N}$ are the valuations of the counters. Let $\text{size}(\gamma) = \nu^1 + \nu^2 + \nu^3$. A configuration $\gamma_2 = (q_2, \nu_2^1, \nu_2^2, \nu_2^3)$ is a lossy successor of a configuration $\gamma_1 = (q_1, \nu_1^1, \nu_1^2, \nu_1^3)$, written $\gamma_1 \rightarrow_M \gamma_2$, if: either (i) $\delta(q) = c_i++$; goto q_2 , $0 \leq \nu_2^i \leq \nu_1^i + 1$ and for all $j \in \{1, 2, 3\} \setminus \{i\}$: $\nu_2^j \leq \nu_1^j$; or (ii) $\delta(q) =$ if $c_i \neq 0$ then c_i-- ; goto q_2 else goto q , $\nu_1^i \neq 0$, $0 \leq \nu_2^i \leq \nu_1^i - 1$ and for all $j \in \{1, 2, 3\} \setminus \{i\}$: $0 \leq \nu_2^j \leq \nu_1^j$; or (iii) $\delta(q) =$ if $c_i \neq 0$ then c_i-- ; goto q else goto q_2 , $\nu_1^i = \nu_2^i = 0$ and for all $j \in \{1, 2, 3\} \setminus \{i\}$: $0 \leq \nu_2^j \leq \nu_1^j$. In particular, for $\delta(q) = \text{halt}$, the configurations with location q have no successor. An infinite run of a 3CM M is an infinite sequence $\rho = \gamma_0, \gamma_1, \dots, \gamma_i, \dots$ of configurations of M such that $\gamma_0 = (q_{in}, 0, 0, 0)$ is the initial configuration and $\gamma_i \rightarrow_M \gamma_{i+1}$ for all $i \geq 0$. We say that ρ is *space-bounded* if there exists $k \in \mathbb{N}$ such that for all $j \geq 0$, $\text{size}(\gamma_j) \leq k$. For a bounded run ρ , we denote the smallest such k by $\text{bound}(\rho)$. We denote by $\text{runs}_B^\infty(M)$ the set of infinite space-bounded runs of M . The *repeated reachability problem* is to decide if $\text{runs}_B^\infty(M) = \emptyset$ for a given 3CM M , and it is undecidable.

Theorem 1 ([15]). *The repeated reachability problem for 3CM is undecidable.*

3 ECL realizability is undecidable

We present a reduction of the repeated reachability problem of 3CM to ECL realizability, showing that the realizability problem for ECL is undecidable. To present our reduction, consider a 3CM $M = \langle c_1, c_2, c_3, Q, q_{in}, \delta \rangle$, and a configuration $\gamma = \langle q, \nu^1, \nu^2, \nu^3 \rangle$ of M . We encode runs and configurations as timed words over the alphabet $\Sigma_{\text{Enc}} = \{a, b_1, b_2, b_3, \text{tick}\} \cup Q$. The configuration γ is encoded as a word of the form $\text{tick } q \ a^{\nu^1} \ b_1 \ a^* \ \text{tick } a^{\nu^2} \ b_2 \ a^* \ \text{tick } a^{\nu^3} \ b_3 \ a^*$ (time stamps omitted). The number of a 's occurring between a tick and the b_i encodes the value of the i th counter (note that the a 's after b_i have no influence on the value of the counters). An infinite bounded run ρ of M is encoded as an infinite sequence of such words, one for each configuration of the run. We require that the total number of a 's in each encoding of a configuration *does not increase* along the run ρ . This requirement is sound since we consider only *bounded runs*. For instance, if we encode the initial configuration by having $\text{bound}(\rho)$ a 's after each b_i , then we are sure to be able to encode the whole run. Moreover, decreasing the total number of a 's can only decrease the counter values which corresponds to the lossy semantics of the machine. Finally, the operations on the counters can be implemented as follows: decrementing (resp. incrementing) counter c_i can be done by switching b_i with the first a on its left (resp. right). If there is no such a , then the counter cannot be decremented (resp. incremented).

We give the conditions that an infinite timed $\theta = (w, \tau)$ word has to satisfy to encode a run $\gamma_0, \gamma_1, \dots, \gamma_i, \dots$ of M . In the sequel, we denote w_i by $w(i)$ and τ_i by $\tau(i)$. The first condition constrains w , the untimed part of θ :

$$\mathbf{C1} \quad w \in (\text{tick} \cdot Q \cdot \mathbf{a}^* \cdot \mathbf{b}_1 \cdot \mathbf{a}^* \cdot \text{tick} \cdot \mathbf{a}^* \cdot \mathbf{b}_2 \cdot \mathbf{a}^* \cdot \text{tick} \cdot \mathbf{a}^* \cdot \mathbf{b}_3 \cdot \mathbf{a}^*)^\omega$$

For $\theta = (w, \tau)$ satisfying **C1**, for $k \geq 0$ and $i \in \{1, 2, 3\}$, let $p_k^{\mathbf{t}^i}$ be the position of the $3k + i$ th **tick** in w and $p_k^{\mathbf{b}_i}$ is the position of the $k + 1$ st \mathbf{b}_i in w . Thus, $p_k^{\mathbf{t}^1}$ is the first position in the encoding of γ_k . Then, **C2** and **C3** constrain the time stamps of the letters:

- C2** The first **tick** appears as the first event: $p_0^{\mathbf{t}^1} = 0$, and a **tick** corresponds to one time unit: for every $k \geq 0$, for $i \in \{1, 2, 3\}$: $\tau(p_k^{\mathbf{t}^i}) = \tau(0) + 3k + (i - 1)$.
- C3** The states of M appear 0 time units after the preceding **tick**: for any $j \geq 0$: $w(j) \in Q$ implies that $\tau(j) = \tau(j - 1)$.

Then, for all $k \geq 0$, the subword of θ with time stamps in the interval $[\tau(0) + 3k, \tau(0) + 3k + 3)$ is of the form **tick** $Q \mathbf{a}^* \mathbf{b}_1 \mathbf{a}^* \text{tick} \mathbf{a}^* \mathbf{b}_2 \mathbf{a}^* \text{tick} \mathbf{a}^* \mathbf{b}_3 \mathbf{a}^*$ and encodes $\gamma_k = (q_k, \nu_k^1, \nu_k^2, \nu_k^3)$ with $q_k = w(p_k^{\mathbf{t}^1} + 1)$, $\nu_k^1 = p_k^{\mathbf{b}_1} - p_k^{\mathbf{t}^1} - 2$, $\nu_k^2 = p_k^{\mathbf{b}_2} - p_k^{\mathbf{t}^2} - 1$ and $\nu_k^3 = p_k^{\mathbf{b}_3} - p_k^{\mathbf{t}^3} - 1$. Thus, θ encodes the infinite sequence $\gamma_0, \gamma_1, \dots, \gamma_i, \dots$ of configurations, yet this sequence is not necessarily a *run* of M , as we need to enforce the semantics of M . This is the purpose of conditions **C4** through **C7** given below. Condition **C4** ensures that the first encoding corresponds to the initial configuration of M . Conditions **C5**, **C6** and **C7** encode the lossy semantics of the machine. In particular, it is important to observe how the relation between two successive values of a given counter, say ν_k^i and ν_{k+1}^i can be encoded as a relation between the *time stamps* of the \mathbf{b}_i 's that appear in the encodings of γ_k and γ_{k+1} . More precisely, conditions **C6** and **C7** ensure that for every $k \geq 1$, every \mathbf{a} or \mathbf{b}_i in the encoding of γ_k is matched by one \mathbf{a} or \mathbf{b}_i exactly three time units before in the encoding of γ_{k-1} . As a consequence, the total number of \mathbf{a} 's does not increase along the run, and the values ν_k^i and ν_{k+1}^i of counter i in two successive configurations can be related by comparing the time stamps of the \mathbf{b}_i 's. For instance, if we want $\nu_{k+1}^i \leq \nu_k^i$, then the \mathbf{b}_i in the $k + 1$ st configuration must appear at most three time units later than the \mathbf{b}_i in the k th configuration, i.e., $\tau(p_{k+1}^{\mathbf{b}_i}) \leq \tau(p_k^{\mathbf{b}_i}) + 3$, and so forth.

- C4** The first portion of the word corresponds to the encoding of the initial configuration of M : $w(0) = \text{tick}$, $w(1) = q_{in}$, and $\nu_0^1 = \nu_0^2 = \nu_0^3 = 0$.
- C5** The time stamps of \mathbf{b}_i are chosen according to the semantics of the machine. For all $k \geq 0$: $\delta(q_k) \neq \text{halt}$ and: (i) $\delta(q_k) = c_i++$; **goto** q' implies that $\tau(p_{k+1}^{\mathbf{b}_i} - 1) \leq \tau(p_k^{\mathbf{b}_i}) + 3$ and $q_{k+1} = q'$; (ii) $\delta(q_k) = \text{if } c_i \neq 0 \text{ then } c_i--$; **goto** q' else **goto** q'' and $\nu_k^i = 0$ implies that $q_{k+1} = q'$ and $\tau(p_{k+1}^{\mathbf{b}_i}) \leq \tau(p_k^{\mathbf{b}_i}) + 3$; and (iii) $\delta(q) = \text{if } c_i \neq 0 \text{ then } c_i--$; **goto** q' else **goto** q'' and $\nu_k^i \neq 0$ implies that $q_{k+1} = q''$ and $\tau(p_{k+1}^{\mathbf{b}_i}) < \tau(p_k^{\mathbf{b}_i}) + 3$.
- C6** All \mathbf{a} 's and \mathbf{b} 's are separated by a strictly positive time delay: for all $j \geq 1$, $w(j) \in \{\mathbf{a}, \mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$ implies that $\tau(j - 1) < \tau(j)$.
- C7** Every \mathbf{a} or \mathbf{b}_i that appears in θ after time stamp $\tau(0) + 3$, i.e., in the encoding of γ_k with $k \geq 1$, is matched by an \mathbf{a} or \mathbf{b}_i exactly three time units before, i.e., in γ_{k-1} . For all $j \geq 1$, if $w(j) \in \{\mathbf{a}, \mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$ and $\tau(j) \geq \tau(0) + 3$, then there exists $i < j$ such that $w(i) \in \{\mathbf{a}, \mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$ and $\tau(i) = \tau(j) - 3$.

It is straightforward to see that a word θ satisfying conditions **C1-C7** encodes a run $\rho_\theta \in \text{runs}_B^\infty(M)$.

Lemma 1. *Let M be a 3CM and θ be an infinite timed word that satisfies **C1-C7**. Then, θ encodes a run $\gamma_0, \gamma_1, \dots, \gamma_i, \dots \in \text{runs}_B^\infty(M)$.*

On the other hand, a run ρ of M can be encoded by a timed word $\text{EncComp}(\rho)$ that satisfies **C1-C7**. Let $\kappa = \text{bound}(\rho)$. For $t \in \mathbb{R}^{\geq 0}$ and $v \in \mathbb{N}$ with $v \leq \kappa$, let $\text{EncVal}(v, \kappa, t) = (\mathbf{a}, t_1) \cdots (\mathbf{a}, t_v) (\mathbf{b}, t_{v+1}) (\mathbf{a}, t_{v+2}) \cdots (\mathbf{a}, t_{\kappa+1})$ where, for any $1 \leq i \leq \kappa+1: t_i = i/(\kappa+2)$. For a configuration $\gamma = (q, \nu^1, \nu^2, \nu^3)$, let $\text{EncConf}(\gamma, b, t) = (\text{tick}, t)(q, t) \cdot \text{EncVal}(\nu^1, \kappa, t) \cdot (\text{tick}, t+1) \cdot \text{EncVal}(\nu^1, \kappa, t+1) \cdot (\text{tick}, t+2) \cdot \text{EncVal}(\nu^3, \kappa, t+2)$. Finally, for $\rho = \gamma_0, \gamma_1, \dots, \gamma_j, \dots \in \text{runs}_B^\infty(M)$, let $\text{EncComp}(\rho)$ be the infinite concatenation of the $\text{EncConf}(\gamma_j, \text{bound}(\rho), 3j)$ for $j \geq 0$.

Lemma 2. *Let M be a 3CM. For all $\rho \in \text{runs}_B^\infty(M)$, the timed word $\text{EncComp}(\rho)$ satisfies **C1-C7**.*

Corollary 1. *Let M be a 3CM. There exists a timed word θ satisfying **C1-C7** if and only if $\text{runs}_B^\infty(M) \neq \emptyset$.*

We have thus reduced the repeated reachability problem for 3CM to the satisfiability of conditions **C1-C7**. Since the satisfiability problem for ECL is decidable, it is not possible to construct an ECL formula whose semantics is equivalent to conditions **C1-C7**. In fact, only **C7** cannot be expressed in ECL. For the other conditions, we propose the Encoding formula given below, where \mathcal{AB} denotes $(\mathbf{a} \vee \mathbf{b}_1 \vee \mathbf{b}_2 \vee \mathbf{b}_3)$, and \mathcal{Q} denotes $(\bigvee_{q \in Q} q)$.

$$\text{Encoding} \equiv \text{tick} \wedge \triangleright_{=0} (q_0 \wedge \bigcirc \mathbf{b}_1) \quad (1)$$

$$\wedge \triangleright_{=1} (\text{tick} \wedge \bigcirc \mathbf{b}_2) \quad (2)$$

$$\wedge \triangleright_{=2} (\text{tick} \wedge \bigcirc \mathbf{b}_3) \quad (3)$$

$$\wedge \square (\text{tick} \rightarrow \triangleright_{=1} \text{tick}) \quad (4)$$

$$\wedge \square (\mathcal{Q} \rightarrow (\ominus \text{tick} \wedge \triangleleft_{=0} \text{tick} \wedge \triangleright_{=3} \mathcal{Q})) \quad (5)$$

$$\wedge \square ((\mathbf{b}_1 \vee \mathbf{b}_2 \vee \mathbf{b}_3) \rightarrow (\neg \mathbf{b}_1 \wedge \neg \mathbf{b}_2 \wedge \neg \mathbf{b}_3) \mathcal{U} \text{tick}) \quad (6)$$

$$\wedge \square (\mathbf{b}_1 \rightarrow (\neg \mathbf{b}_1 \wedge \neg \mathbf{b}_3) \mathcal{U} \mathbf{b}_2) \quad (7)$$

$$\wedge \square (\mathbf{b}_2 \rightarrow (\neg \mathbf{b}_1 \wedge \neg \mathbf{b}_2) \mathcal{U} \mathbf{b}_3) \quad (8)$$

$$\wedge \square (\mathbf{b}_3 \rightarrow (\neg \mathbf{b}_2 \wedge \neg \mathbf{b}_3) \mathcal{U} \mathbf{b}_1) \quad (9)$$

$$\wedge \square ((\mathcal{AB} \vee \mathcal{Q} \vee \text{tick}) \rightarrow \triangleright_{>0} (\mathcal{AB} \vee \text{tick})) \quad (10)$$

$$\wedge \bigwedge_{q \in Q} \square \text{instr}(q) \quad (11)$$

where, for $q \in Q$, the formula $\text{instr}(q)$ is defined as follows:

1. If $\delta(q) = i++$; goto q' , then:

$$\text{instr}(q) \equiv q \rightarrow \triangleright_{=3} q' \wedge \left(\left(\text{inc}_i \wedge \bigwedge_{j \neq i} \text{keep}_j \right) \vee \left(\bigwedge_j \text{keep}_j \right) \right) \quad (12)$$

2. If $\delta(q) = \text{if } i \neq 0 \text{ then } i--; \text{ goto } q' \text{ else goto } q''$, then:

$$\text{instr}(q) \equiv (q \wedge \text{isnull}_i) \rightarrow \left(\triangleright_{=3} q'' \wedge \bigwedge_j \text{keep}_j \right) \quad (13)$$

$$\wedge (q \wedge \neg \text{isnull}_i) \rightarrow \left(\triangleright_{=3} q' \wedge \text{dec}_i \wedge \bigwedge_{j \neq i} \text{keep}_j \right) \quad (14)$$

3. If $\delta(q) = \text{halt}$, then:

$$\text{instr}(q) \equiv \square \left(\bigwedge_{q \in Q} \neg q \right) \quad (15)$$

The formulas dec_i , inc_i and keep_i are defined as follows. For $i \in \{1, 2, 3\}$: $\text{dec}_i \equiv \triangleright_{<3} (\mathbf{b}_i \wedge \triangleright_{<3} \mathbf{b}_i)$; $\text{inc}_i \equiv \triangleright_{<3} (\mathbf{b}_i \wedge \bigcirc \mathbf{a} \wedge \triangleright_{\leq 3} (\mathbf{a} \wedge \bigcirc \mathbf{b}_i))$; and $\text{keep}_i \equiv \triangleright_{<3} (\mathbf{b}_i \wedge \triangleright_{\leq 3} \mathbf{b}_i)$. Finally, $\text{isnull}_1 \equiv (\neg \mathbf{a}) \mathcal{U} \mathbf{b}_1$ and for $i = 2, 3$: $\text{isnull}_i \equiv \triangleright_{=i-1} ((\neg \mathbf{a}) \mathcal{U} \mathbf{b}_i)$. It is easy to see that Encoding corresponds to conditions C1-C6:

Lemma 3. *For all timed words θ , $\theta \in \llbracket \text{Encoding} \rrbracket$ if and only if θ satisfies C1-C6.*

Corollary 2. *Let M be a 3CM. There exists a timed word $\theta \in \llbracket \text{Encoding} \rrbracket$ satisfying C7 if and only if $\text{runs}_B^\infty(M) \neq \emptyset$.*

To conclude the proof that ECL realizability is undecidable, we show how timed games can be exploited to check whether there exists a timed word θ that satisfies Encoding and C7, and hence whether $\text{runs}_B^\infty(M) \neq \emptyset$. The game we consider is $G_M = \langle \Sigma_{\text{Enc}}, \{c\}, \llbracket \varphi_M \rrbracket \rangle$, and we show that Player 1 has a winning strategy in G_M iff $\text{runs}_B^\infty(M) \neq \emptyset$. Before we formally define φ_M , we give some intuition. In this game, we use the winning condition φ_M to force Player 1 to faithfully simulate M by satisfying conditions C1-C7. Note that Player 1 controls the full alphabet of the configuration's encoding. However, by Lemma 3 defining $\varphi_M = \text{Encoding}$ is not sufficient: Player 1 could *cheat* by inserting extra a's in the play, in order to increase the values of the counters. We use the game interaction with Player 2 to force condition C7. Using action c , Player 2 will be given the possibility to check that Player 1 does not increase the counters as follows. First, Player 2 is allowed to play at most one c , and only exactly 0 time unit after an \mathbf{a} or a \mathbf{b}_i . In this case, we say that Player 2 performs a *check*, and the meaning of this c is to pinpoint a particular \mathbf{a} or \mathbf{b}_i in the word that should correspond to a previous \mathbf{a} or \mathbf{b}_i three time units before, as stated in C7. If it is not the case, then we say that Player 2 has *detected an error*, and thus C7 is violated. Hence, the second ingredient is to let Player 1 loose whenever an *error is detected*, i.e. when a c appears right after an \mathbf{a} or a \mathbf{b}_i that is not preceded exactly three time units before by a corresponding \mathbf{a} or \mathbf{b}_i .

These constraints on the number and positions of the c 's and on the detection of the errors turn out to be expressible in ECL. By combining these constraints with Encoding, we obtain $\varphi_M \equiv \text{Hyp} \rightarrow \text{Goal}$ where:

$$\text{Hyp} \equiv \square \left(c \rightarrow (\triangleleft_{=0} \mathcal{AB}) \right) \wedge \left((\neg c \wedge \triangleright_{\geq 3} c) \vee \square \neg c \right) \wedge \square (c \rightarrow \square(\neg c))$$

ensures that Player 2 performs the checks right after an a or a b_i has been produced, not in the first configuration, and at most once. Moreover we let $\text{Goal} \equiv \text{Encoding} \wedge \text{Check}$, with $\text{Check} \equiv \diamond c \rightarrow \diamond (\mathcal{AB} \wedge \triangleright_{=3} c)$. Goal ensures that Player 1 generates a word that satisfies conditions C1-C6, and that she loses whenever she cheats: whenever she plays an a or a b_i that is not preceded by a corresponding a or b_i exactly three time units before, Player 2 can play a c (provided that she hasn't played a c before) that will falsify Check , and thus φ_M .

Let us show there is a winning strategy in $G_M = \langle \Sigma_{\text{Enc}}, \{c\}, \llbracket \text{Hyp} \rightarrow \text{Goal} \rrbracket \rangle$ for Player 1 iff $\text{runs}_B^\infty(M) \neq \emptyset$. The 'if' direction is easy, since Player 1 can play according to $\text{EncComp}(\rho)$, for any $\rho \in \text{runs}_B^\infty(M)$. Indeed, since $\text{EncComp}(\rho)$ satisfies condition C7, Player 2 will never detect an error.

Proposition 1. *Let M be a 3CM. If $\text{runs}_B^\infty(M) \neq \emptyset$, then Player 1 has a winning strategy in the timed game $G_M = \langle \Sigma_{\text{Enc}}, \{c\}, \llbracket \text{Hyp} \rightarrow \text{Goal} \rrbracket \rangle$.*

Let us finally show that, if Player 1 has a winning strategy, then $\text{runs}_B^\infty(M) \neq \emptyset$. The idea of the proof is as follows. We first observe that, by definition of φ_M , Player 1 can win the game if Player 2 does not satisfy Goal or if she decides to check an a or a b_i which is preceded by an a or a b_i three time units before (then Player 2 cannot make further checks, which leaves to Player 1 the ability to cheat in the rest of the play). In this case Player 1 wins without having to faithfully simulate M . Of course, Player 2 has a better strategy to choose the action c exactly 0 time unit after the first wrong a or b_i has been issued. Since Player 1 has to win when Player 2 plays in this way, a winning strategy for Player 1 has to ensure that Goal holds, i.e. that Encoding and C7 are satisfied, thus faithfully simulating an infinite run of M .

In other words, we consider a strategy StratEnc for Player 2 that forces Player 1 to play according to Encoding and C7. Given a strategy π_1 for Player 1, define StratEnc as follows: for every finite prefix $\theta = (w, \tau)$ of length ℓ , let $\text{StratEnc}(\theta) = (c, 0)$ (i.e., Player 2 is detecting an error) if and only if (i) $w(\ell - 1) \in \{a, b_1, b_2, b_3\}$, (ii) $\tau(\ell - 1) \geq 3$, (iii) there is no $k < \ell - 1$ such that $\tau(\ell - 1) - \tau(k) = 3$ and $w(k) \in \{a, b_1, b_2, b_3\}$, and (iv) there is no $k' < \ell$ such that $w(k') = c$; otherwise, we let $\text{StratEnc}(\theta) = (c, \Delta + 1)$ where Δ is the time delay proposed by Player 1 when she plays according to π_1 , i.e., $\pi_1(\theta) = (\alpha, \Delta)$ for some $\alpha \in \Sigma_{\text{Enc}}$. The next lemma says that, against strategy StratEnc for Player 2, a winning strategy of Player 1 produces a play satisfying Goal and C7.

Lemma 4. *Let π_1 be a winning strategy for Player 1 in G_M . Then, for all plays $\theta \in \text{Outcome}(G_M, \pi_1, \text{StratEnc})$: $\theta \models \text{Goal}$ and θ satisfies C7.*

Proposition 2. *Let M be a 3CM. If Player 1 has a winning strategy in G_M , then $\text{runs}_B^\infty(M) \neq \emptyset$.*

By Theorem 1 and Proposition 1 and 2 we obtain the following result.

Theorem 2. *The realizability problem for ECL is undecidable.*

It is easy to extend this undecidability result to the td-realizability problem for ECL, since the winning strategy presented in the proof of Proposition 1 is also winning for $\text{WC}_1(\llbracket \psi \rrbracket)$, and against strategy StratEnc for Player 2, a winning strategy of Player 1 for $\text{WC}_1(\llbracket \psi \rrbracket)$ also produces plays that satisfy C7.

Theorem 3. *The td-realizability problem for ECL is undecidable.*




Fig. 1. A deterministic parity timed automaton for φ_e . States 1, 3, 4 have priority 0, and states 2, 5 have priority 1.

4 Positive result on LTL_{\triangleleft}

In this section we show that the realizability problem is decidable for the syntactic fragment LTL_{\triangleleft} . More precisely, we present an algorithm to solve the realizability and td-realizability problems for LTL_{\triangleleft} .

Given a timed game $G = \langle \Sigma_1, \Sigma_2, \llbracket \psi \rrbracket \rangle$ for $\psi \in LTL_{\triangleleft}$, the main idea of the algorithm consists in building a *deterministic timed automaton with parity condition* D_ψ that accepts exactly the winning words for Player 1, i.e., $L(D_\psi) = WC_1(\llbracket \psi \rrbracket)$. This automaton can then be used to build a winning strategy for Player 1 (if it exists), using the techniques of [9].

First observe that we do not need to remember the exact time stamps of every event in a timed word to evaluate the truth value of a formula ψ of LTL_{\triangleleft} . Indeed, there are only finitely many subformulas of the form $\triangleleft_I \varphi$ in ψ , and these are the only real-time formulas. Intuitively, we can thus consider ψ as an LTL formula over the augmented alphabet $\Sigma \times 2^P$ where P is a set of proposition that tracks the truth values of the $\triangleleft_I \varphi$ subformulas. Such untimed words are called *Hintikka sequences* of ψ , and we first show that we can build a nondeterministic finite automaton A_ψ with generalized Büchi condition that accepts those Hintikka sequences. After determinization of A_ψ (giving B_ψ), we translate B_ψ into a deterministic timed automaton C_ψ with parity condition, by relating the truth value of the propositions that track the subformulas $\triangleleft_I \varphi$ with the value of clocks of the automaton. We get $L(C_\psi) = \llbracket \psi \rrbracket$. For td-realizability, we use the construction of [9] to construct a deterministic timed automaton D_ψ that accounts the time-diverging condition on timed words. Thus, $L(D_\psi) = WC_1(\llbracket \psi \rrbracket)$. The automaton D_ψ can be used to extract a td-winning strategy for Player 1, and the automaton C_ψ to extract a winning strategy for Player 1 [9]. The automaton C_{φ_e} of the formula φ_e of Example 2 is given in Fig. 1.

An *Hintikka sequence* of a formula ψ is an (untimed) word h over the alphabet $\Sigma \times 2^P$ where $P = \{p_\varphi \mid \varphi \equiv \triangleleft_I \varphi_1 \text{ is a subformula of } \psi\}$. The semantics of h is the set $\llbracket h \rrbracket$ of timed words (w, τ) over Σ such that $h(i) = (w(i), \Omega_i)$ where $\Omega_i = \{p_\varphi \in P \mid (w, \tau), i \models \varphi\}$, for all $i \geq 0$. Note that for all Hintikka sequences $h \neq h'$, we have $\llbracket h \rrbracket \cap \llbracket h' \rrbracket = \emptyset$. Therefore, given a timed word $\theta = (w, \tau)$, we denote by $Hs(\theta)$ the unique Hintikka sequence h such that $\theta \in \llbracket h \rrbracket$, and given a language L of timed words, we denote by $Hs(L)$ the set $\{Hs(\theta) \mid \theta \in L\}$.

Lemma 5. For all LTL_{\triangleleft} formula ψ , we have $\llbracket Hs(\llbracket \psi \rrbracket) \rrbracket = \llbracket \psi \rrbracket$.

Given an $\text{LTL}_{\triangleleft}$ formula ψ , we denote by $\text{Sub}(\psi)$ the set containing the formulas φ and $\diamond\varphi$ for all subformulas φ of ψ . From ψ , we construct the following nondeterministic (untimed) finite automaton with generalized Büchi condition on edges, $A_\psi = \langle Q, q_{in}, E, \alpha \rangle$ over the alphabet $\Sigma \times 2^P$ ($P = \{p_\varphi \mid \varphi \equiv \triangleleft_1 \varphi_1 \text{ is a subformula of } \psi\}$).

- Q contains q_{in} and all the $q \subseteq \text{Sub}(\psi)$ that are *consistent*. A subset q is consistent iff: (i) there exists a unique $a \in \Sigma$ such that $a \in q$; (ii) for all subformulas φ_1, φ_2 of ψ , if $\varphi_2 \equiv \neg\varphi_1$, then $\varphi_1 \in q$ iff $\varphi_2 \notin q$; and (iii) for all subformulas $\varphi_1 \vee \varphi_2$ of ψ , $\varphi_1 \vee \varphi_2 \in q$ iff $\varphi_1 \in q$ or $\varphi_2 \in q$.
- $E \subseteq Q \times (\Sigma \times 2^P) \times Q$ contains all edges (q, σ, q') such that $\sigma = (a, \{p_\varphi \in P \mid \varphi \in q'\})$ where $\{a\} = \Sigma \cap q'$ and, either (i) $q = q_0$, $\psi \in q'$ and $\varphi_1 \mathcal{S} \varphi_2 \notin q'$ for all formulas $\varphi_1 \mathcal{S} \varphi_2 \in \text{Sub}(\psi)$, or (ii) $q \neq q_0$, for all subformula $\varphi_1 \mathcal{U} \varphi_2$ of ψ , we have $\varphi_1 \mathcal{U} \varphi_2 \in q$ iff either (a) $\varphi_2 \in q'$, or (b) $\varphi_1 \in q'$ and $\varphi_1 \mathcal{U} \varphi_2 \in q'$; and for all subformula $\varphi_1 \mathcal{S} \varphi_2$ of ψ , we have $\varphi_1 \mathcal{S} \varphi_2 \in q'$ iff either (a) $\varphi_2 \in q$, or (b) $\varphi_1 \in q$ and $\varphi_1 \mathcal{S} \varphi_2 \in q$.
- α is a set of accepting sets of edges, containing for each subformula $\varphi_1 \mathcal{U} \varphi_2$ of ψ the set $\{(q, \sigma, q') \in E \mid \varphi_1 \mathcal{U} \varphi_2 \notin q \text{ or } \varphi_2 \in q'\}$.

Lemma 6. *For all $\text{LTL}_{\triangleleft}$ formula ψ , we have $\llbracket L(A_\psi) \rrbracket = \llbracket \psi \rrbracket$.*

The next lemma is crucial to translate B_ψ (the deterministic version of A_ψ) into a timed automaton C_ψ . Indeed, in the time automaton C_ψ , we use one clock for each formula of the form $\triangleleft_1 \varphi$ to remember the last time φ has been true. Lemma 7 shows that only the information about the past of the word is relevant to know when these clocks have to be reset.

Lemma 7. *For all nonempty (untimed) finite words w over the set of propositions Σ , for all runs r_1, r_2 of A_ψ over w , the states $\text{Last}(r_1)$ and $\text{Last}(r_2)$ contain exactly the same PastECL formulas.*

From A_ψ , we obtain a deterministic (untimed) automaton B_ψ with parity condition such that $L(B_\psi) = L(A_\psi)$ by Piterman's determinization procedure [21]. The states of B_ψ are Safra trees s , whose root $\text{root}(s)$ tracks the standard subset construction. Therefore, by Lemma 7, for every transition (s, σ, s') of B_ψ , all states $q \in \text{roots}'$ agree on the PastECL subformulas of ψ . So, we can define a (deterministic) timed automaton C_ψ over alphabet Σ and clocks $\{x_\varphi \mid \triangleleft_1 \varphi \text{ is a subformula of } \psi\}$ as follows: the state space of C_ψ is a copy of the state space of B_ψ , and for each transition $(s, (a, \Omega), s')$ in B_ψ , if for all $p_\varphi \in \Omega$ with $\varphi \equiv \triangleleft_1 \varphi_1$, we have $\diamond\varphi_1 \in \text{root}(s)$, then there is a transition (s, g, a, R, s') in C_ψ such that: $R = \{x_\varphi \mid p_\varphi \in \Omega\}$ and g is the conjunction of (i) all constraints $x_{\varphi_1} \in I$ s.t. $p_\varphi \in \Omega$ and $\varphi \equiv \triangleleft_1 \varphi_1$ is a subformula of ψ , and (ii) all constraints $x_\varphi \notin I$ s.t. $p_\varphi \notin \Omega$, $\varphi \equiv \triangleleft_1 \varphi_1$ is a subformula of ψ , and $\diamond\varphi_1 \in \text{root}(s)$.

Proposition 3. *For all PastECL formula ψ , the timed automaton C_ψ with parity condition is deterministic and $L(C_\psi) = \llbracket \psi \rrbracket$.*

Using the results of [9], a deterministic timed automaton D_ψ with parity condition can be constructed from C_ψ such that $L(D_\psi) = \text{WC}_1(\llbracket \psi \rrbracket)$. The number of locations of D is $O(|C| \cdot d)$ where d is the number of priorities in C_ψ , and the number of priorities in D is $d + 2$. To decide if Player 1 has a winning strategy for $\text{WC}_1(\llbracket \psi \rrbracket)$, we evaluate a μ -calculus fixpoint formula [10] that computes the set of winning states of Player

1 for the winning condition $WC_1(\llbracket\psi\rrbracket)$. The μ -calculus formula uses a *controllable predecessor operator* $CPre(Z)$ that computes the set of states in which Player 1 can force the game to Z in one move. The controllable predecessor operator preserves the regions of the timed automaton D_ψ , i.e., if Z is a union of regions, $CPre(Z)$ is also a union of regions. Therefore, the winning states of Player 1 can be computed in time $O(|D_\psi| \cdot m! \cdot 2^m \cdot (2c+1)^m)^d$ where $|D_\psi|$ is the number of locations in D_ψ , m is the number of clocks, c is the largest constant, and d is the maximal priority in D_ψ [9]. If we let $n = |A_\psi|$, we get $d = 2 + 2 \cdot n \cdot O(|\psi|)$, $c = c_\psi$ is the largest constant that occurs as an integer endpoint of an interval I in a subformula $\triangleleft_I \varphi$ of ψ , m is the number of subformula $\triangleleft_I \varphi$ of ψ , and $|D_\psi| = 2d \cdot n^n \cdot n!$ [9, 21]. This is at most $2^{O(2^{O(|\psi|)})} \cdot (2c_\psi + 1)^{2^{O(|\psi|)}}$ where $|\psi|$ is the length of ψ .

Theorem 4. *For $\psi \in LTL_{\triangleleft}$ over alphabet $\Sigma_1 \uplus \Sigma_2$, deciding whether Player 1 is *td-winning* the game $\langle \Sigma_1, \Sigma_2, \llbracket\psi\rrbracket \rangle$ can be done in time $2^{O(2^{O(|\psi|)})} \cdot (2c_\psi + 1)^{2^{O(|\psi|)}}$.*

The realizability problem for LTL_{\triangleleft} can be solved by the same technique as in Theorem 4, using the automaton C_ψ instead of D_ψ .

Theorem 5. *For $\psi \in LTL_{\triangleleft}$ over alphabet $\Sigma_1 \uplus \Sigma_2$, deciding whether Player 1 is *winning* the game $\langle \Sigma_1, \Sigma_2, \llbracket\psi\rrbracket \rangle$ can be done in time $2^{O(2^{O(|\psi|)})} \cdot (2c_\psi + 1)^{2^{O(|\psi|)}}$.*

Since the realizability problem for LTL is 2EXPTIME-hard, we get the following corollary.

Corollary 3. *The realizability and *td-realizability* problems for LTL_{\triangleleft} are 2EXPTIME-complete.*

5 Discussion

We close the paper by mentioning several open problems for future works. First, several semantical models have been proposed for real-time behaviors [4]. We conjecture that our proofs of (un)decidability extend to the case where the real-time models are timed state sequences, i.e. finite variable functions from $\mathbb{R}^{\geq 0}$ to Σ , and that our decidability result extends to LTL with the past formulas of MITL (the intuition is that the formulas of past MITL can be translated to deterministic timed automata [16]). Second, the realizability problem for ECL remains open in the case of finite words (the reachability problem is decidable for 3CM). It is our belief that techniques based on well-quasi orderings [20] should be investigated. Then, one could consider restricted classes of strategies (such as strategies with imperfect information [8], or with bounded resources [6]) to recover decidability. Finally, our positive result relies on the Safra construction for determinization. A Safraless procedure [14, 11] should be investigated.

References

1. R. Alur and D.L. Dill. A Theory of Timed Automata. *TCS*, 126(2), 1994.
2. R. Alur, T. Feder, and T. Henzinger. The benefits of relaxing punctuality. *J. ACM*, 43(1), 1996.

3. R. Alur, L. Fix, and T. Henzinger. Event-clock automata: a determinizable class of timed automata. *TCS*, 211(1-2), 1999.
4. R. Alur and T. Henzinger. Logics and models of real time: A survey. *Proc. REX Workshop*, 1992. Springer.
5. R. Alur and T. Henzinger. A really temporal logic. *J. ACM*, 41(1), 1994.
6. P. Bouyer, L. Bozzelli, and F. Chevalier. Controller synthesis for MTL specifications. *Proc CONCUR'06*, LNCS 4137, 2006, Springer.
7. F. Cassez, A. David, E. Fleury, K. G. Larsen, and D. Lime. Efficient on-the-fly algorithms for the analysis of timed games. *Proc CONCUR'05*, LNCS 3653, 2005, Springer.
8. F. Cassez, A. David, E. Fleury, K. G. Larsen, D. Lime and J.F. Raskin Timed Control with Observation Based and Stuttering Invariant Strategies. *Proc ATVA'07*, LNCS 4762, 2005, Springer.
9. L. de Alfaro, M. Faella, T. Henzinger, R. Majumdar, and M. Stoelinga. The element of surprise in timed games. *Proc CONCUR'03*, LNCS 2761, 2003, Springer.
10. L. de Alfaro, T. Henzinger, and R. Majumdar. From verification to control: Dynamic programs for omega-regular objectives. *Proc. LICS'01*, IEEE Computer Society Press, 2001.
11. E. Filliot, N. Jin and J.F. Raskin An Antichain Algorithm for LTL Realizability. *Proc. CAV'09*, to appear.
12. T. Henzinger, J.-F. Raskin and P.-Y. Schobbens. The Regular Real-Time Languages *Proc. ICALP'98*, LNCS 1443, 1998, Springer.
13. R. Koymans. Specifying real-time properties with metric temporal logic. *RT Syst.*, 2(4), 1990.
14. O. Kupferman and M. Vardi. Safraless decision procedures. *Proc. FOCS'05*, 2005, IEEE Computer Society.
15. R. Mayr. Undecidable problems in unreliable computations. *TCS*, 297(1-3), 2003.
16. O. Maler, D. Nickovic, and A. Pnueli. Real time temporal logic: Past, present, future. *Proc. FORMATS'05*, LNCS 3829, 2005, Springer.
17. O. Maler, D. Nickovic, and A. Pnueli. On synthesizing controllers from bounded-response properties. *Proc. CAV'07*, LNCS 4590, 2007, Springer.
18. Z. Manna and A. Pnueli. *Temporal verification of reactive systems: safety*. 1995, Springer.
19. O. Maler, A. Pnueli, and J. Sifakis. On the synthesis of discrete controllers for timed systems. *Proc STACS'95*, LNCS 900, 1995, Springer.
20. J. Ouaknine and J. Worrell. On the decidability of metric temporal logic. In *Proc LICS '05* IEEE Computer Society Press, 2005.
21. N. Piterman. From nondeterministic Büchi and Streett automata to deterministic parity automata. *LMCS*, 3(3), 2007.
22. A. Pnueli. The temporal logic of programs. *Proc. SFCS'77*, 1977, IEEE Computer Society.
23. A. Pnueli and R. Rosner. On the synthesis of a reactive module. *Proc. POPL'89*, 1989, ACM.
24. J.-F. Raskin. *Logics, Automata and Classical Theories for Deciding Real Time*. PhD thesis, FUNDP (Belgium), 1999.
25. J.-F. Raskin and P.-Y. Schobbens. The logic of event clocks: decidability, complexity and expressiveness. *Automatica*, 34(3), 1998.
26. P. Wolper. The tableau method for temporal logic: An overview. *Logique et Analyse*, (110–111), 1985.

A Proofs of the lemmata

Lemma 1. *Let M be a 3CM and θ be an infinite timed word that satisfies C1-C7.*

Then, θ encodes a run $\gamma_0, \gamma_1, \dots, \gamma_i, \dots \in \text{runs}_B^\infty(M)$.

Proof. By induction on k .

Lemma 2. *Let M be a 3CM. For all $\rho \in \text{runs}_B^\infty(M)$, the timed word $\text{EncComp}(\rho)$ satisfies C1-C7.*

Proof. By careful inspection of the definitions.

Lemma 3. *For all timed words θ , $\theta \in \llbracket \text{Encoding} \rrbracket$ if and only if θ satisfies C1-C6.*

Proof (sketch). Constraints (1-9) correspond to conditions C1-C4. Condition C5 corresponds to (11). Condition C6 corresponds to (10). \square

Lemma 4. *Let π_1 be a winning strategy for Player 1 in G_M . Then, for all plays $\theta \in \text{Outcome}(G_M, \pi_1, \text{StratEnc})$: $\theta \models \text{Goal}$ and θ satisfies C7.*

Proof. Let us first show that $\theta \models \text{Goal}$. Since π_1 is a winning strategy, the play $\theta = (w, \tau)$ is winning and $\theta \models \text{Hyp} \rightarrow \text{Goal}$. Hence, it suffices to show that $\theta \models \text{Hyp}$. First, observe that by definition of *StratEnc*, if there is p^c such that $w(p^c) = c$, then $\tau(p^c - 1) = \tau(p^c) \geq 3$ and $w(p^c - 1) \in \{a, b_1, b_2, b_3\}$. Then it suffices to show that there is at most one $p^c \in \mathbb{N}$ such that $w(p^c) = c$. Let p^c be the first position in θ where Player 2 plays $(c, 0)$ and plays first. In this case, clearly $w(p^c) = c$. Moreover, for every $k > p^c$, Player 1 will play (α, Δ) , for some $\alpha \in \Sigma_1$ and $\Delta \in \mathbb{R}^{\geq 0}$, according to her strategy π_1 , and Player 2 will play $(c, \Delta + 1)$, according to *StratEnc*. Hence, Player 2 will never play first again starting from $p^c + 1$, and thus, for every $k > p^c + 1$, $w(k) \neq c$. Therefore $\theta \models \text{Hyp}$.

Now, let us show that *StratEnc* also guarantees that $\theta = (w, \tau)$ satisfies C7. This is done *by contradiction*. Assume that $\theta \not\models \text{C7}$, and assume that $\theta = (\alpha_0, \tau_0)(\alpha_1, \tau_1) \dots$. Let p be the first position such that $w(p) \in \{a, b_1, b_2, b_3\}$, $\tau(p) \geq 3$ and there is no $p' < p$ with $w(p') \in \{a, b_1, b_2, b_3\}$ and $\tau(p') = \tau(p) - 3$. Then, Player 2 detects and error and performs a check by playing $(c, 0)$. As a consequence, there is in $\text{Outcome}(G_M, \pi_1, \text{StratEnc})$ a play θ' with prefix $(\alpha_0, \tau_0)(\alpha_1, \tau_1) \dots (\alpha_p, \tau_p)(c, \tau_p)$. However, $\theta' \not\models \text{Goal}$ since Checkis not satisfied, thus π_1 is not a winning strategy. Contradiction.

Lemma 5. *For all LTL $_{\triangleleft}$ formula ψ , we have $\llbracket \text{Hs}(\llbracket \psi \rrbracket) \rrbracket = \llbracket \psi \rrbracket$.*

Proof (sketch). The inclusion $\llbracket \psi \rrbracket \subseteq \llbracket \text{Hs}(\llbracket \psi \rrbracket) \rrbracket$ directly follows from the definitions. To establish that $\llbracket \text{Hs}(\llbracket \psi \rrbracket) \rrbracket \subseteq \llbracket \psi \rrbracket$, consider a timed word θ such that $\text{Hs}(\theta) \in \text{Hs}(\llbracket \psi \rrbracket)$, i.e., $\text{Hs}(\theta) = \text{Hs}(\theta')$ for some $\theta' \models \psi$. It is easy to see that $\theta \models \psi$ by viewing ψ as standard (untimed) LTL formula over the alphabet $\Sigma \times 2^P$ whose truth value is determined by the Hintikka sequence. \square

Lemma 6. *For all LTL $_{\triangleleft}$ formula ψ , we have $\llbracket L(A_\psi) \rrbracket = \llbracket \psi \rrbracket$.*

Proof (sketch). The automaton A_ψ essentially follows the tableaux construction of [26, 18] for LTL formulas over the alphabet $\Sigma \times 2^P$. Therefore, by Lemma 5, for all Hintikka sequences $h \in L(A_\psi)$ and all timed words $\theta \in \llbracket h \rrbracket$, we have $\theta \in \llbracket \psi \rrbracket$. On the other hand, if $\theta \in \llbracket \psi \rrbracket$, then $\text{Hs}(\theta)$ satisfies ψ when viewed as an LTL formula over $\Sigma \times 2^P$, and thus $\text{Hs}(\theta) \in L(A_\psi)$. \square

Lemma 7. *For all nonempty (untimed) finite words w over the set of propositions Σ , for all runs r_1, r_2 of A_ψ over w , the states $\text{Last}(r_1)$ and $\text{Last}(r_2)$ contain exactly the same PastECL formulas.*

Proof (sketch). By induction on the length of w . For words of length 1 (i.e., $w = w_0$), consider a run $r = q_0 w_0 q_1$ of A_ψ . Note that the state q_1 contains no formula of the form $\psi_1 \mathcal{S} \psi_2$, and therefore the PastECL formulas in q_1 are Boolean combinations of symbols in $\Sigma \times 2^P$. By definition of A_ψ , the letter $w_0 \in \Sigma \times 2^P$ uniquely determines which action in Σ and which formulas of the form $\triangleleft_1 \varphi$ are in q_1 . Moreover, by definition of the state space Q of A_ψ , this uniquely determines which Boolean combinations of subformulas are in q_1 .

The induction step is left to the reader (for each transition (q, w_i, q') of A_ψ , which PastECL formulas are in q' is uniquely determined by the subformulas of q and the letter w_i). \square

B Proofs of the propositions

Proposition 1. *Let M be a 3CM. If $\text{runs}_B^\infty(M) \neq \emptyset$, then Player 1 has a winning strategy in the timed game $G_M = \langle \Sigma_{\text{Enc}}, \{c\}, \llbracket \text{Hyp} \rightarrow \text{Goal} \rrbracket \rangle$.*

Proof. To prove the property, we show how to build a winning strategy π_1 for Player 1, under the assumption that there exists a run ρ in $\text{runs}_B^\infty(M)$. We consider the word $\text{EncComp}(\rho) = (w, \tau)$, and define the strategy π_1 as follows. For any prefix $\theta' = (w', \tau')$ of length ℓ , we let $\pi_1(\theta') = (w(p), \tau(p) - \tau'(\ell - 1))$ where:

$$p = \max \left\{ L \mid \left(\begin{array}{l} \exists h = \{0, \dots, L-1\} \mapsto \{0, \dots, \ell-1\} : \\ \forall 0 \leq j < L-1 : h(j) < h(j+1) \text{ and} \\ \forall 0 \leq j < L : w(j) = w'(h(j)) \wedge \tau(j) = \tau'(h(j)) \end{array} \right) \right\}$$

That is, we look for the maximal prefix of $\text{EncComp}(\rho)$ that appears as a subword⁵ of θ' , and propose to play the first unseen letter of $\text{EncComp}(\rho)$. Remark that the set whose maximum is taken is never empty, because $i = 0$ always satisfies the constraint. Thus, this definition is sound.

Let us show that π_1 is winning, i.e., for all $\theta = (w, \tau) \in \text{Outcome}_1(G_M, \pi_1)$, we have $\theta \models \varphi_M$. First observe that, in each of the following cases, $\theta \models \varphi_M$ because $\theta \not\models \text{Hyp}$:

1. If there are two positions $p_1 \neq p_2$ such that $w(p_1) = w(p_2) = c$.
2. If there is a position p such that $w(p) = c$ and $\tau(p) < 3$.

⁵ We need to consider subwords since Player 2 could have inserted a c , and Player 1 has to keep playing consistently with $\text{EncComp}(\rho)$, no matter what Player 2 does.

3. If there is a position $p > 0$ such that $w(p) = c$ and $w(p-1) \notin \{a, b_1, b_2, b_3\}$.
4. If there is a position $p > 0$ such that $w(p) = c$ and $\tau(p) \neq \tau(p-1)$.

Thus, let us assume that θ falsifies all of the above, hence $\theta \models \text{Hyp}$. As a consequence, θ is a word obtained by inserting at most one c exactly 0 time unit right after an a or a b_i in $\text{EncComp}(\rho)$. More precisely: either $\theta = w_1 \cdot (\alpha, t) \cdot (c, t) \cdot w_2$, with $w_1 \cdot (\alpha, t) \cdot w_2 = \text{EncComp}(\rho)$ and $\alpha \in \{a, b_1, b_2, b_3\}$, or $\theta = \text{EncComp}(\rho)$. In the latter case, $\theta \models \text{Goal}$, by Lemma 2. In the former case, it is routine to check that $\theta \models \text{Goal}$ too because each a or b_i that appears after time stamp $\tau(0) + 3$ in $\text{EncComp}(\rho)$ corresponds to an a or a b_i three time units before, and thus Check, which is the only constraint on the c 's appearing in Goal, is satisfied too. Thus, $\theta \models \varphi_M$. \square

Proposition 2. *Let M be a 3CM. If Player 1 has a winning strategy in G_M , then $\text{runs}_B^\infty(M) \neq \emptyset$.*

Proof. Let π_1 be a winning strategy for Player 1 and let us consider a play θ from $\text{Outcome}(G_M, \pi_1, \text{StratEnc})$. By Lemma 4, $\theta \models \text{Goal}$ and thus $\theta \in \llbracket \text{Encoding} \rrbracket$. Moreover, by Lemma 4 again, θ satisfies C7. Hence, $\text{runs}_B^\infty(M) \neq \emptyset$ by Corollary 2. \square

Proposition 3. *For all PastECL formula ψ , the timed automaton C_ψ with parity condition is deterministic and $L(C_\psi) = \llbracket \psi \rrbracket$.*

Proof (sketch). A key property of the tableaux construction A_ψ is that for all runs $r = q_0 w_0 q_1 w_1 \dots$ of A_ψ on w , for all formulas $\psi_1 \in \text{Sub}(\psi)$ and for all $i \geq 0$, we have $w, i \models \psi_1$ if and only if $\psi_1 \in q_{i+1}$ (where ψ_1 is viewed as an LTL formula over the alphabet $\Sigma \times 2^P$ by replacing formulas $\triangleleft_I \varphi$ by propositions $p_{\triangleleft_I \varphi} \in P$, for which the satisfiability relation \models is independent of the time sequence) [26, 18]. Let $\varphi \equiv \triangleleft_I \varphi_1$ be a subformula of ψ . Using Lemma 7, and since x_{φ_1} is reset on every transition from s to s' with label containing p_φ (or equivalently such that φ occurs in every state $q' \in \text{roots}'$ by the above), it is easy to show by induction that for all runs $r = (q_0, v_0)(w_0, \tau_0)e_0(q_1, v_1)(w_1, \tau_1)e_1 \dots$ of C_ψ over $\theta = (w, \tau)$, for all $i \geq 1$, if there exists $j \leq i$ such that $\theta, j \models \varphi$ and $\theta, k \not\models \varphi$ for all $k, j < k \leq i$, then $v_i(x_{\varphi_1}) = \tau(i) - \tau(j)$. Now, the construction of the transition guard g of a transition (s, g, a, R, s') in C_ψ corresponding to $(s, (a, \Omega), s')$ in B_ψ ensures that on all such runs r , we have $\theta, j \models \triangleleft_I \varphi$ if and only if $p_{(\triangleleft_I \varphi)} \in \Omega$. Therefore, $\theta \in L(C_\psi)$ if and only if $\text{Hs}(\theta) \in L(A_\psi)$. That $L(C_\psi) = \llbracket \psi \rrbracket$ then follows from Lemma 6.

The proof that C_ψ is deterministic is left to the reader. \square