You May Delay, but Time Will Not⋆: Timed Games Under Delayed Control⋆⋆

Kim G. Larsen[1][0000−0002−5953−3384] and Martin Zimmermann[1][0000−0002−8038−2453]

Aalborg University, Aalborg, Denmark
{kgl,mzi}@cs.aau.dk

Abstract. Inspired by Martin Fränzle's persistent and influential work on capturing and handling delay inherent to cyber-physical systems in the formal verification of such systems, we study timed games where controllable actions do not take effect immediately, but only after some delay, i.e., they are scheduled for later execution.

We show that solving such games is undecidable if an unbounded number of actions can be pending. On the other hand, we present a doubly-exponential time algorithm for games with a bound on the number of pending actions, based on a reduction to classical timed games. This makes timed games under delayed control with bounded schedules solvable with existing tools like UPPAAL.

Keywords: Timed Games · Delayed Control · Synthesis.

1 Introduction

For more than 30 years, Martin Fränzle has been a main contributor to the foundations of real-time and cyber-physical systems with contributions to modeling, specification, monitoring, shielding, model checking and synthesis. Some of the earliest work of Martin Fränzle was targeting the expressive and very elegant Duration Calculus [7] with contributions to efficient model checking [11] and even synthesis [10]. The latter work was presented by Martin Fränzle at the conference FTRTFT already in 1996 at Uppsala University coinciding with the very first tutorial on UPPAAL [4] provided by the first author of this paper. In all likelihood, the conference FTRTFT marked the first encounter of the two.

The most recent collaboration between Martin Fränzle and (now) both authors of this contribution was concerned with capturing delays in the monitoring of real-time systems. Such delays are an inherent feature of all complex cyber-physical systems, but are often ignored due to the complexity they incur. However, delays may be of crucial importance for the faithful analysis of a given system. They may arise, for example, in the following situations:

⋆ "Poor" Richard Saunders, otherwise known as Benjamin Franklin.

Processing and Computation Delays: Many control systems use complex algorithms, like object recognition in autonomous vehicles or predictive models in industrial control. These computations take time, creating delays before control commands are issued. To ensure accurate control, data from sensors often needs to be filtered to reduce noise, which introduces additional processing time. Also, digital control systems operate on discrete samples, so a low sampling rate can lead to delays in recognizing system changes and responding to them.

Actuator Dynamics: Actuators, like motors and hydraulic systems, have physical limitations in how quickly they can respond. For example, the time it takes for a large motor to accelerate or decelerate causes delays. Furthermore, systems with large masses or physical resistance, like robotic arms or ship rudders, experience delays because they must overcome inertia before moving.

Communication and Transmission Delays: In networked control systems, especially those connected remotely, data transmission introduces delays. For instance, a remote control signal sent to a drone might take time to reach it over the internet. For long-distance systems, such as satellite controls, there is an inherent delay because of the time it takes for signals to travel.

Synchronization and Coordination: Some systems require synchronization involving multiple components or operations. For example, in manufacturing, a robotic arm might need to wait for a previous operation to finish before starting the next. In complex systems, certain tasks are prioritized, which can lead to delays in lower-priority control actions, especially in multi-tasking environments like autonomous vehicles or robots.

Intentional Delay for Stability and Safety: Control actions may be delayed by design to ensure conditions are safe for operation. For example, chemical plants often have safety verifications before initiating changes. Furthermore, to prevent oscillations or over-corrections, some systems use intentional delays, such as damping or time lags, to smooth out responses and ensure stability.

All of the above factors can lead to delays between when a control command is issued and when the action actually occurs, impacting the responsiveness of the system and often requiring compensatory design measures. The importance of delays in control systems was very early on recognized by Martin Fränzle[1] and, in a series of groundbreaking results, he made contributions to the theoretical foundations of real-time and cyber-physical systems which *include delays as a first-class citizen*. We recall in particular:

Delay Differential Equations: In several papers, Martin Fränzle and collaborators have studied and applied Delay Differential Equations (DDEs) as a natural model of networked control systems, where the communication delay in the feedback loop cannot always be ignored. Iterating bounded degree

---

[1] In relation to the title of this paper, one may say that Martin Fränzle did *not* delay when it came to dealing with delay! Rather, he acted before many of the rest of us.

interval-based Taylor over-approximations of the time-wise segments of the solution to a DDE has been proposed as a method for analysing stability and safety [25]. Further contributions on DDE's by Martin Fränzle include the verification of temporal logic [22, 23].

Safety Shielding under Delay: Shields are correct-by-construction runtime enforcers that guarantee safe execution by correcting any action that may cause a violation of a formal safety specification. Most recently, Martin Fränzle together with co-authors proposed synthesis algorithms to compute delay-resilient shields that guarantee safety under worst-case assumptions on the delays of the input signals [9]. This work also introduced novel heuristics for deciding between multiple corrective actions, designed to minimize the number of future shield interferences caused by delays.

Controller Synthesis under Delay: Martin Fränzle and his co-authors studied the problem of synthesizing controllers that can effectively operate in the presence of delays during interactions with their environments [8]. The authors model this problem using two-player games, where one player represents the controller and the other represents the environment. The paper explores the complexity of finding winning strategies for the controller under various types of delay, including constant delays, out-of-order delivery of messages, and bounded message loss. In particular a novel incremental algorithm for synthesizing delay-resilient controllers is presented that outperforms existing reduction-based approaches [24].

More recently, Martin Fränzle, Sarah Winter, and the second author of this paper [13, 14] have studied the relation between controller synthesis under delay as discussed above and so-called delay games, another model of controller synthesis under delay due to Hosch and Landweber [15], thereby allowing to transfer results for delay games [17, 18] to the model introduced by Martin Fränzle and his co-authors.

Monitoring under Delay: The two authors of this paper – together with Thomas M. Grosen – have had the pleasure of a recent collaboration with Martin Fränzle on online monitoring of real-time systems in the setting of delayed observations of actions [12]. In fact, a purely zone-based online monitoring algorithm has been presented for MITL specifications, which handles *parametric* delays without recurrence to costly verification procedures for parametric timed automata. The monitoring algorithm has been implemented on top of the real-time model checking tool UPPAAL [19] with encouraging initial results.

The joint work on monitoring was initiated during a research visit of Martin Fränzle in Aalborg in early 2023. Here, one of the research questions on the joint wish-list was precisely *delayed control* in a timed game setting, thereby generalizing the work on controller synthesis from the discrete-time setting to the real-time setting. We are happy to provide an initial study of this question with this Festschrift contribution.

To provide a first illustration of the concepts, Figure 1 shows a LEGO Mindstorm production system involving a box moving on a conveyor belt. Now the
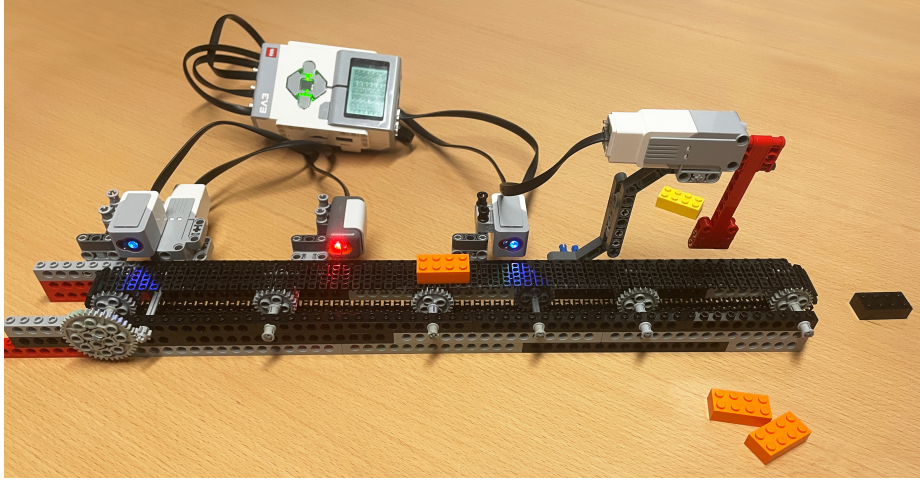
Fig. 1. LEGO Mindstorm Production System.

timed automata $\mathcal{P}$ and $\mathcal{S}$ in Figure 2 model the production system. The various locations of the timed automaton $\mathcal{P}$ indicate the position and potential treatment of the box, e.g. in Hole a hole is being drilled into the box, and in Paint the box is being painted (obviously, not all these treatments are implemented in the LEGO model). In the first four locations the box may at any point in time be subject to the control action kick (indicated by full transitions). In the first three locations (On, Hole, Paint) this will lead to an error-state (Err_i), whereas if kick'ed in the Piston-location, the box will correctly end in the Succ-location. If the box is not kick'ed it will erroneously end in Off. The invariants and guards on the clock x of the timed automaton $\mathcal{P}$ ensure that the time spent in each (top-) location is between 8 and 10 seconds. The uncontrollability of the exact time, when the box mov'es from one location to the next, is indicated by the use of dashed transitions.

Now, the challenge is to make sure that the control action kick is issued (only) when the box is in the Piston location. Cassez et al. considered this problem in the setting where the controller has *partial information* of the position of the box in the system [6]. Here, we are in the setting of control actions being delayed, i.e., there is some (known) delay from when a control action is scheduled until it actually is realized. The timed automaton $\mathcal{S}$ models the possibilities of scheduling kick-actions. In particular, $\mathcal{S}$ has two scheduling actions, schd(kick, 13) respectively schd(kick, 22), labeling the dotted transitions of $\mathcal{S}$. Taking these transitions schedules kick after a delay of 13 respectively 22 seconds. The local clock y of $\mathcal{S}$ ensures that kick actions are scheduled with a minimum time-separation of 8 seconds. Now the synthesis problem is to find a strategy for when schd(kick, 13) respectively schd(kick, 22) are scheduled, so that the box is appropriately kick-ed in the Piston location to reach the goal location Succ. It can be deduced that when arriving to the Hole location, the box will definitely
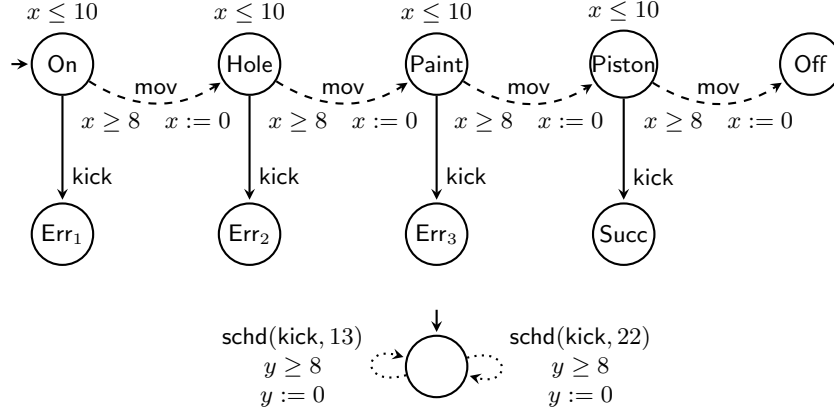
Fig. 2. The production system $\mathcal{P}$ (top) and the scheduler $\mathcal{S}$ (bottom).

be in the Piston location within 20 to 24 seconds. Thus, scheduling $\mathsf{schd}(\mathsf{kick}, 22)$ precisely when arriving at the Hole location provides an adequate strategy. In contrast, an increased timing uncertainty where the residence time of the box in each location is between 6 and 10 seconds implies that there is no strategy ensuring that the goal location is reached.

In the following, we introduce timed games under delayed control as informally described above. Note that in our setting, actions are scheduled and take effect after some delay. This is in fact opposed to the setting considered by Martin Fränzle and his co-authors when they studied the discrete case of delayed control [8]. There, actions take effect immediately, but perception of the current state the system is in is delayed. As argued by Martin Fränzle and his co-authors, these two settings are actually equivalent. For technical reasons, we prefer here, without loss of generality, to take the perspective of *actions taking effect with delay* instead of *perception being delayed*.

Our main contributions are as follows: We show that solving timed games under delayed control is in general undecidable, even for reachability and safety conditions and without uncontrollable actions (i.e., in a one-player setting). The undecidability crucially relies on an unbounded number of actions being pending. If we disallow this and impose that only a bounded number of actions can be pending, then we regain decidability. In fact, we present an algorithm with doubly-exponential running time based on a reduction to classical timed games (without delay). Finally, we show that our generalization of timed games to timed games under delayed control is conservative in the sense that each timed game (without delay) can be turned into an equivalent timed game under delayed control. This implies that the ExpTime-hardness of solving timed games also applies to solving timed games under delayed control.

2 Preliminaries

The set of nonnegative integers is denoted by $\mathbb{N}$, the set of integers by $\mathbb{Z}$, and the set of nonnegative reals by $\mathbb{R}_{\geq 0}$.

In the following subsections, we introduce timed automata [1] and timed games [2, 20], mostly following the notation presented by Cassez et al. in [5].

2.1 Clocks, Clock Constraints, and Zones

Let X be a finite set of (real-valued) clocks. A (clock) valuation of X is a mapping $v\colon X \to \mathbb{R}_{\geq 0}$. We write $\mathbf{0}$ for the valuation mapping every clock to 0. For $X' \subseteq X$, we write $v[X']$ for the valuation mapping every clock $x \in X'$ to 0 and every clock $x \notin X'$ to $v(x)$, i.e., the clocks in X' are reset. For $\delta \in \mathbb{R}_{\geq 0}$, we write $v + \delta$ for the valuation mapping each clock x to $v(x) + \delta$, i.e., δ units of time pass.

The set $\mathcal{C}(X)$ of clock constraints over X is defined by the grammar

$$\varphi ::= x \sim k \mid x - y \sim k \mid \varphi \wedge \varphi$$

where x, y range over X and k ranges over $\mathbb{Z}$. Let $\mathcal{B}(X)$ denote the clock constraints over X that do not use atomic constraints of the form $x - y \sim k$, i.e., $\mathcal{B}(X)$ contains conjunctions of atomic constraints of the form $x \sim k$.

For a clock constraint $\varphi \in \mathcal{C}(X)$ and a valuation v of X, we write $v \models \varphi$ if v satisfies φ, which is defined as expected. We write $[\![\varphi]\!]$ for the set of valuations that satisfy φ. A zone Z is a set of valuations such that $Z = [\![\varphi]\!]$ for some $\varphi \in \mathcal{C}(X)$.

2.2 Timed Automata

A timed automaton $\mathcal{A} = (L, \ell_0, \mathsf{Act}, X, E, \mathsf{Inv})$ consists of a finite set L of locations containing the initial location ℓ_0, a finite set Act of actions, a finite set X of clocks, a finite set $E \subseteq L \times \mathcal{B}(X) \times \mathsf{Act} \times 2^X \times L$ of edges, and a function $\mathsf{Inv}\colon L \to \mathcal{B}(X)$ assigning invariants to locations. As usual, we measure the size of a timed automaton in the number of locations, the number of clocks, and the largest constant appearing in a guard or invariant, as these factors alone determine the complexity of zone-based algorithms, which we rely on here.

A state of $\mathcal{A}$ is a pair (ℓ, v) where $\ell \in L$ and v is a valuation of X, the initial state is $(\ell_0, \mathbf{0})$. In a state, either time can pass (as long as the state's invariant is satisfied) or an edge is taken: intuitively, (ℓ, g, a, X', ℓ') leads from ℓ to ℓ', is labeled by the action a, can be executed if the guard g is satisfied by the current valuation, and the new valuation is obtained by resetting the clocks in X' and has to satisfy the invariant of ℓ'. Formally, $\mathcal{A}$ induces the labeled transition system $\mathfrak{T}(\mathcal{A})$ whose set of vertices is the set of states of $\mathcal{A}$, whose initial vertex is the initial state of $\mathcal{A}$, and where we have two types of transitions:

- Time transitions: For $\delta \in \mathbb{R}_{\geq 0}$, there is a transition $(\ell, v) \xrightarrow{\delta} (\ell, v')$ if $v' = v + \delta$ and $v + \delta' \models \mathsf{Inv}(\ell)$ for all $0 \leq \delta' \leq \delta$.

– Discrete transitions: For $a \in \mathsf{Act}$, there is a transition $(\ell, v) \xrightarrow{a} (\ell', v')$ if there is an edge $(\ell, g, a, X', \ell') \in E$ such that $v \models g$, $v' = v[X']$, and $v' \models \mathsf{Inv}(\ell')$.

A run of $\mathcal{A}$ is a (finite or infinite) alternating sequence of time and discrete transitions in $\mathfrak{T}(\mathcal{A})$. If ρ is a finite run, we write $\mathsf{Last}(\rho)$ for the last state of ρ. An infinite run ρ is time-divergent if the sum of the δ labeling time transitions in ρ is infinite.

Example 1. The sequence

$$(\mathsf{On}, x = 0) \xrightarrow{8.5} (\mathsf{On}, x = 8.5) \xrightarrow{\mathsf{mov}} (\mathsf{Hole}, x = 0) \xrightarrow{9.1}$$

$$(\mathsf{Hole}, x = 9.1) \xrightarrow{\mathsf{mov}} (\mathsf{Paint}, x = 0) \xrightarrow{9.5} (\mathsf{Paint}, x = 9.5) \xrightarrow{\mathsf{mov}}$$

$$(\mathsf{Piston}, x = 0) \xrightarrow{8.1} (\mathsf{Piston}, x = 8.1) \xrightarrow{\mathsf{kick}} (\mathsf{Succ}, x = 8.1)$$

is a run of the timed automaton $\mathcal{P}$ of Figure 2.

2.3   Timed Games

A timed game $\mathcal{G} = (\mathcal{A}, \mathsf{Act}_c, \mathsf{Act}_u, F)$ consists of a timed automaton $\mathcal{A}$ whose set of actions is the disjoint union of the set Act_c of controllable actions and the set Act_u of uncontrollable actions, and where F is a subset of $\mathcal{A}'s$ locations, used to define winning conditions.

A strategy for $\mathcal{G}$ is a partial function f mapping finite runs in $\mathfrak{T}(\mathcal{A})$ to $\mathsf{Act}_c \cup \{\lambda\}$ such that for every finite run ρ with $f(\rho) \neq \lambda$ there is a state s such that $\mathsf{Last}(\rho) \xrightarrow{f(\rho)} s$. We say that f is state-based if $\mathsf{Last}(\rho) = \mathsf{Last}(\rho')$ implies $f(\rho) = f(\rho')$ for all finite runs ρ, ρ'.

The set $\mathsf{Outcomes}(s, f)$ of outcomes of f from a state s of $\mathcal{A}$ is defined inductively as follows:

– $s \in \mathsf{Outcomes}(s, f)$, and
– if a finite run ρ is in $\mathsf{Outcomes}(s, f)$ and $\mathsf{Last}(\rho) \xrightarrow{a} s'$ is a transition in $\mathfrak{T}(\mathcal{G})$, then $\rho \xrightarrow{a} s'$ is in $\mathsf{Outcomes}(s, f)$ if one of the following conditions holds:
  - $a \in \mathsf{Act}_u$,
  - $a \in \mathsf{Act}_c$ and $a = f(\rho)$, or
  - $a \in \mathbb{R}_{\geq 0}$ and for all $0 \leq \delta < a$ there exists a state s'' of $\mathcal{G}$ such that $\rho \xrightarrow{\delta} s''$ and $f(\rho \xrightarrow{\delta} s'') = \lambda$.
– An infinite run ρ is in $\mathsf{Outcomes}(s, f)$ if all of its finite prefixes are in $\mathsf{Outcomes}(s, f)$.

A run ρ is maximal if it is infinite, or if it is finite and $\mathsf{Last}(\rho)$ does not have any successors in $\mathfrak{T}(\mathcal{A})$. Let s_0 denote the initial state of $\mathcal{A}$. A strategy f

– satisfies reachability if every maximal run in $\mathsf{Outcomes}(s_0, f)$ visits a location in F at least once,
– satisfies safety if every (finite or infinite) run in $\mathsf{Outcomes}(s_0, f)$ visits only locations in F, and

– satisfies time-divergence if every maximal run is infinite and time-divergent.

Proposition 1 ([16]). *The following problem is* EXPTIME*-complete: Given a timed game, is there a strategy that satisfies reachability?*

Automatic synthesis of real-time controllers for timed games was first considered by Asarin, Maler, Pnueli and Sifakis [2, 20], who proved decidability for both safety and reachability conditions. Also, they showed that state-based strategies suffice in both cases. An efficient zone-based on-the-fly synthesis algorithm for timed games was presented by Cassez et al. [5] and is available in the tool UPPAAL Tiga [3].

Example 2. Reconsider the timed game $\mathcal{P}$ from Figure 2 with $\mathsf{Act}_c = \{\mathsf{kick}\}$, $\mathsf{Act}_u = \{\mathsf{mov}\}$, and $F = \{\mathsf{Succ}\}$. Given the objective to reach F, it can be seen that the following partial function f_{Succ} is a winning (state-based) strategy (note that we use a zone-based description):

$$f_{\mathsf{Succ}}(\mathsf{On}, 0 \leq x \leq 10) = \lambda \qquad f_{\mathsf{Succ}}(\mathsf{Hole}, 0 \leq x \leq 10) = \lambda$$
$$f_{\mathsf{Succ}}(\mathsf{Paint}, 0 \leq x \leq 10) = \lambda \qquad f_{\mathsf{Succ}}(\mathsf{Piston}, 0 \leq x < 8) = \mathsf{kick}$$

3   Timed Games under Delayed Control

In this section, we introduce our model of timed games under delayed control, which is a generalization of timed games, and give an example. The following sections are then devoted to the study of this model.

In a classical timed game, when a controllable action $a \in \mathsf{Act}_c$ is selected by a strategy, then an edge labeled by a is taken immediately to continue the run. In timed games under delayed control, when a strategy selects an action $a \in \mathsf{Act}_c$, then this comes with some associated delay t and an edge labeled by a is only taken after t units of time have passed. Thus, a strategy intuitively schedules a for execution in t units of time. Thus, between the selection of an action and the actual execution, other (previously scheduled) controllable actions and uncontrollable actions may be executed, which can change the state of the game. To reiterate, a strategy in a timed game under delayed control selects pairs (a, t) of actions and delays, not actions. Thus, we have three types of actions now:

– uncontrollable actions,
– scheduling actions of the form $\mathsf{schd}(a, t)$ where $a \in \mathsf{Act}_c$ and t is a delay, and
– actions in Act_c.

As the actions in Act_c are now no longer under direct control by a strategy, we call them *control* actions instead of controllable actions. The controllable actions, i.e., those that can be selected by a strategy, are the scheduling actions.

Formally, a timed game under delayed control $\mathcal{D} = (\mathcal{A}, \mathsf{Act}_c, \mathsf{Act}_u, T, F)$ consists of a timed automaton $\mathcal{A}$ whose set of actions is the disjoint union of

– the set Act_c of control actions,

- the set Act_u of uncontrollable actions, and
- the set $\mathsf{Schd}(\mathsf{Act}_c, T) = \{\mathsf{schd}(a, t) \mid a \in \mathsf{Act}_c, t \in T\}$ of scheduling actions,

where $T \subseteq \mathbb{N}$ is a finite set of natural numbers, and where F is a subset of $\mathcal{A}$'s locations, used to define winning conditions.

We require the following form of deadlock freedom: For each location ℓ of $\mathcal{A}$, every clock valuation v, and every control action $a \in \mathsf{Act}_c$, there is an edge (ℓ, g, a, X', ℓ') in $\mathcal{A}$ such that $v \models g$ and $v[X'] \models \mathsf{Inv}(\ell')$. This implies that a previously scheduled action can be executed in every possible state.

Remark 1. This form of deadlock freedom can always be achieved by adding a fresh sink state and fresh edges that are enabled exactly for those states and control actions for which no original edge is enabled. When the fresh sink is not in F, then the resulting timed game under delayed control is equivalent to the original one.

A schedule σ is a sequence $(a_0, t_0) \cdots (a_{k-1}, t_{k-1}) \in (\mathsf{Act}_c \times \mathbb{R}_{\geq 0})^*$ such that $t_0 \leq t_1 \leq \cdots \leq t_{k-1}$. A state of $\mathcal{D}$ is a triple (ℓ, v, σ), where (ℓ, v) is a state of $\mathcal{A}$ and σ is a schedule. The initial state is $(\ell_0, \mathbf{0}, \varepsilon)$, where $(\ell_0, \mathbf{0})$ is the initial state of $\mathcal{A}$, i.e., the schedule is initially empty. Intuitively, a schedule contains the scheduled control actions a_i and timestamps t_i, which the intuition that a_i has to be executed in t_i units of time. Thus, when time passes, say δ units, then the timestamps are decreased by δ.

Now, in a state of $\mathcal{D}$, time can pass, an edge with an uncontrollable action can be taken, a control action a can be scheduled for execution in t units of time (action $\mathsf{schd}(a, t)$), or a scheduled action can be executed (if its timestamp is zero). Formally, $\mathcal{D}$ induces the labeled transition system $\mathfrak{T}(\mathcal{D})$ whose set of vertices is the set of states of $\mathcal{D}$, whose initial vertex is the initial state of $\mathcal{D}$, and where we have four types of transitions:

- Time transitions: For $\delta \in \mathbb{R}_{\geq 0}$, there is a transition

$$(\ell, v, (a_0, t_0) \cdots (a_{k-1}, t_{k-1})) \xrightarrow{\delta} (\ell, v', (a_0, t_0 - \delta) \cdots (a_{k-1}, t_{k-1} - \delta))$$

  if $v' = v + \delta$, $\delta \leq t_0$, and $v + \delta' \models \mathsf{Inv}(\ell)$ for all $0 \leq \delta' \leq \delta$. Note that the timestamps in the schedule all decrease by δ.
- Uncontrolled discrete transitions: For $a \in \mathsf{Act}_u$, there is a transition

$$(\ell, v, \sigma) \xrightarrow{a} (\ell', v', \sigma)$$

  if there is an edge $(\ell, g, a, X', \ell') \in E$ such that $v \models g$, $v' = v[X']$, and $v' \models \mathsf{Inv}(\ell')$.
- Scheduled discrete (control) transitions: For $a \in \mathsf{Act}_c$, there is a transition

$$(\ell, v, (a_0, t_0) \cdots (a_{k-1}, t_{k-1})) \xrightarrow{a} (\ell', v', (a_1, t_1) \cdots (a_{k-1}, t_{k-1}))$$

  if $a = a_0$, $t_0 = 0$, and there is an edge $(\ell, g, a, X', \ell') \in E$ such that $v \models g$, $v' = v[X']$, and $v' \models \mathsf{Inv}(\ell')$. Due to the deadlock freedom we have imposed, each vertex of the form $(\ell, v, (a_0, t_0) \cdots (a_{k-1}, t_{k-1}))$ with $t_0 = 0$ has at least one outgoing transition labeled by a_0, which removes (a_0, t_0) from the schedule.

– Scheduling discrete (control) transitions: For $a \in \mathsf{Act}_c$ and $t \in T$, there is a transition
$$(\ell, v, (a_0, t_0) \cdots (a_{k-1}, t_{k-1})) \xrightarrow{\mathsf{schd}(a,t)}$$
$$(\ell', v', (a_0, t_0) \cdots (a_{i-1}, t_{i-1})(a, t)(a_i, t_i) \cdots (a_{k-1}, t_{k-1})$$

if $t_{i-1} \leq t \leq t_i$ (where $t_{-1} = 0$) and there is an edge $(\ell, g, \mathsf{schd}(a,t), X', \ell') \in E$ such that $v \models g$, $v' = v[X']$, and $v' \models \mathsf{Inv}(\ell')$. Note that there is an (uncontrollable) degree of freedom in where (a, t) is added to the schedule if there is (at least) one pair (a_j, t_j) with $t_j = t$. In that situation, (a, t) can be added before or after (a_j, t_j) and transitions corresponding to both choices are in $\mathfrak{T}(\mathcal{D})$. In general, if there are k' pairs with timestamp $t_j = t$, then there are $k' + 1$ places where (a, t) can be added. Said differently, actions are ordered when they are inserted in the schedule and then executed in that order, but we take all (valid) orderings into account.

In a timed game under delayed control, control actions have to be scheduled. Once they are scheduled, they will be executed after their specified delay has passed. Hence, a strategy in a timed game under delayed control only schedules control actions. Formally, a strategy f for $\mathcal{D}$ is a partial function mapping finite runs in $\mathfrak{T}(\mathcal{D})$ to $\mathsf{Schd}(\mathsf{Act}_c, T) \cup \{\lambda\}$ such that for every finite run ρ with $f(\rho) \neq \lambda$ there is a state s such that $\mathsf{Last}(\rho) \xrightarrow{f(\rho)} s$.

The set $\mathsf{Outcomes}(s, f)$ of outcomes of f from a state s of $\mathcal{D}$ is defined inductively as follows:

– $s \in \mathsf{Outcomes}(s, f)$, and
– if a finite run ρ is in $\mathsf{Outcomes}(s, f)$ and $\mathsf{Last}(\rho) \xrightarrow{a} s'$ is a transition in $\mathfrak{T}(\mathcal{D})$, then $\rho \xrightarrow{a} s'$ is in $\mathsf{Outcomes}(s, f)$ if one of the following conditions holds:
  - $a \in \mathsf{Act}_c \cup \mathsf{Act}_u$ (recall that control actions are scheduled and then executed after their delay has passed),
  - $a \in \mathsf{Schd}(\mathsf{Act}_c, T)$ and $a = f(\rho)$, or
  - $a \in \mathbb{R}_{\geq 0}$ and for all $0 \leq \delta < a$ there exists a state s'' of $\mathcal{D}$ such that $\rho \xrightarrow{\delta} s''$ and $f(\rho \xrightarrow{\delta} s'') = \lambda$.
– An infinite run ρ is in $\mathsf{Outcomes}(s, f)$ if all of its finite prefixes are in $\mathsf{Outcomes}(s, f)$.

A run ρ is maximal if it is infinite, or if it is finite and $\mathsf{Last}(\rho)$ does not have any successors in $\mathfrak{T}(\mathcal{D})$. Let s_0 denote the initial state of $\mathcal{D}$. A strategy f

– satisfies reachability if every maximal run in $\mathsf{Outcomes}(s_0, f)$ visits a location in F at least once,
– satisfies safety if every (finite or infinite) run in $\mathsf{Outcomes}(s_0, f)$ visits only locations in F, and
– satisfies time-divergence if every maximal run is infinite and time-divergent.

Example 3. The product of $\mathcal{P}$ and $\mathcal{S}$ from Figure 2 constitutes a timed game under delayed control, where kick-actions can be scheduled with a delay of 13

or 22 with a minimum time separation of 8 seconds. The objective of reaching Succ is satisfied by, e.g., one of the following two strategies f_{Succ}^{22} and f_{Succ}^{13}:

$$f_{\mathsf{Succ}}^{22}(\mathsf{On}, 0 \leq x \leq 10, \varepsilon) = \lambda$$
$$f_{\mathsf{Succ}}^{22}(\mathsf{Hole}, 0 \leq x < 2 \wedge y \geq 8, \varepsilon) = \mathsf{schd}(\mathsf{kick}, 22)$$
$$f_{\mathsf{Succ}}^{13}(\mathsf{On}, 0 \leq x \leq 10, \varepsilon) = \lambda$$
$$f_{\mathsf{Succ}}^{13}(\mathsf{Hole}, 0 \leq x \leq 10, \varepsilon) = \lambda$$
$$f_{\mathsf{Succ}}^{13}(\mathsf{Paint}, 0 \leq x < 3 \wedge y \geq 8, \varepsilon) = \mathsf{schd}(\mathsf{kick}, 13)$$

In the following, we study the problem of determining whether, for a given timed game under delayed control, there is a strategy satisfying a fixed winning condition from the list above. In Section 4, we prove that for unbounded schedules, the problem is undecidable both for reachability and safety conditions. Thus, we turn our attention to bounded schedules in Section 5: This restriction is sufficient to obtain decidability (in fact membership in 2ExpTime). This result is shown by modeling a timed game under delayed control with bounded schedules as a timed game. Finally, in Section 6, we prove that timed games under delayed control are a conservative extension of timed games, i.e., each timed game can be modeled by a timed game under delayed control.

4   Undecidability for Unbounded Schedules

In this section, we show that determining whether a given timed game under delayed control has a strategy that satisfies reachability (safety) is undecidable. This result relies on the unboundedness of the schedule. It is well-known that finite automata with an unbounded queue have an undecidable emptiness problem, as one can use the queue to simulate the tape of a Turing machine. Here, we use similar ideas, but prefer to simulate two-counter machines using timed games under delayed control. Just as problems for automata with a queue are already undecidable, we construct our timed games under delayed control without uncontrollable actions, i.e., undecidability stems from the queue, not the fact that we consider a game.

A two-counter machine $\mathcal{M}$ is a sequence

$$(0 : \mathtt{I}_0)(1 : \mathtt{I}_1) \cdots (k - 2 : \mathtt{I}_{k-2})(k - 1 : \mathtt{STOP}),$$

where the first element of a pair $(\ell : \mathtt{I}_\ell)$ is the line number and $\mathtt{I}_\ell$ for $\ell \in \{0, \ldots, k - 2\}$ is an instruction of the form

- $\mathtt{INC}(\mathtt{X}_i)$ with $i \in \{0, 1\}$,
- $\mathtt{DEC}(\mathtt{X}_i)$ with $i \in \{0, 1\}$, or
- $\mathtt{IF}\ \mathtt{X}_i\mathtt{=0}\ \mathtt{GOTO}\ \ell'\ \mathtt{ELSE}\ \mathtt{GOTO}\ \ell''$ with $i \in \{0, 1\}$ and $\ell', \ell'' \in \{0, \ldots, k - 1\}$.

A configuration of $\mathcal{M}$ is of the form (ℓ, c_0, c_1) with $\ell \in \{0, \ldots, k-1\}$ (the current line number) and $c_0, c_1 \in \mathbb{N}$ (the current contents of the counters $\mathtt{X}_0$ and $\mathtt{X}_1$, respectively). The initial configuration is $(0, 0, 0)$ and the unique successor configuration of a configuration (ℓ, c_0, c_1) is defined as follows:

- If $\mathtt{I}_\ell = \mathtt{INC}(\mathtt{X}_i)$, then the successor configuration is $(\ell + 1, c_0', c_1')$ with $c_i' = c_i + 1$ and $c_{1-i}' = c_{1-i}$.
- If $\mathtt{I}_\ell = \mathtt{DEC}(\mathtt{X}_i)$, then the successor configuration is $(\ell + 1, c_0', c_1')$ with $c_i' = \max\{c_i - 1, 0\}$ and $c_{1-i}' = c_{1-i}$.
- If $\mathtt{I}_\ell = \mathtt{IF}\ \mathtt{X}_i\mathtt{=0}\ \mathtt{GOTO}\ \ell'\ \mathtt{ELSE}\ \mathtt{GOTO}\ \ell''$ and $c_i = 0$, then the successor configuration is (ℓ', c_0, c_1).
- If $\mathtt{I}_\ell = \mathtt{IF}\ \mathtt{X}_i\mathtt{=0}\ \mathtt{GOTO}\ \ell'\ \mathtt{ELSE}\ \mathtt{GOTO}\ \ell''$ and $c_i > 0$, then the successor configuration is (ℓ'', c_0, c_1).
- If $\mathtt{I}_\ell = \mathtt{STOP}$, then (ℓ, c_0, c_1) has no successor configuration.

The unique run of $\mathcal{M}$ (starting in the initial configuration) is defined as expected. It is either finite (line $k - 1$ is reached) or infinite (line $k - 1$ is never reached). In the former case, we say that $\mathcal{M}$ terminates.

Proposition 2 ([21]). *The following problem is undecidable: Given a two-counter machine $\mathcal{M}$, does $\mathcal{M}$ terminate?*

Next, we show that termination of two-counter machines reduces to the existence of strategies satisfying reachability, which implies that the latter problem is undecidable as well. In the proof, we make crucial use of the fact that the schedule may be unbounded.

Theorem 1. *The following problem is undecidable: Given a timed game under delayed control, is there a strategy that satisfies reachability?*

Proof. Given a two-counter machine

$$\mathcal{M} = (0 : \mathtt{I}_0)(1 : \mathtt{I}_1) \cdots (k - 2 : \mathtt{I}_{k-2})(k - 1 : \mathtt{STOP}),$$

we construct a timed game under delayed control $\mathcal{D}_\mathcal{M}$ that has a strategy that satisfies reachability if and only if $\mathcal{M}$ terminates.

The timed automaton of $\mathcal{D}_\mathcal{M}$ has locations $0, 1, \ldots, k - 1$ as well as some auxiliary locations (typically named m, possibly with decorations), the initial location ℓ_0, and a sink location ℓ_S. The set F of target locations contains only the location $k - 1$, which is a sink as well. The timed automaton uses a single clock x and the set of control actions is $\mathsf{Act}_c = \{Z, O, \#\}$ while the set of uncontrolled actions is empty. Finally, the set T of delays with which an action can be scheduled is the singleton $\{1\}$.

A schedule σ encodes a valuation (c_0, c_1) of the two counters $\mathtt{X}_0$ and $\mathtt{X}_1$ where c_0 (c_1) is the number of Z's (O's) in σ, i.e., the action $\#$ is irrelevant here. Later, it is used as a divider to implement the correct update of the counters. Thus, configurations of $\mathcal{M}$ are encoded by pairs of locations of $\mathcal{D}_\mathcal{M}$ and schedules. Note that the clock valuation is irrelevant for the encoding.

In our construction, we only consider schedules $\sigma = (a_0, t_0) \cdots (a_{k-1}, t_{k-1})$ satisfying certain properties. We say that σ is well-timed if we have $0 < t_0 < t_1 < \cdots t_{k-1} \leq 1$, i.e., the first action cannot be executed immediately, no two actions are scheduled with the same timestamp, and all of them have to be executed before one unit of time has passed. We use the clock to enforce well-timedness. Furthermore, we say that σ is well-sorted if $k \geq 1$, i.e., σ is nonempty, and $a_j \in \{Z, O\}$ for all $j < k-1$ and $a_{k-1} = \#$, i.e., exactly the last scheduled action is $\#$.

In the following, we explain how to simulate the computation of $\mathcal{M}$ by $\mathcal{D}_{\mathcal{M}}$ using gadgets implementing the three types of instructions. As already alluded to, we use the schedule to encode valuations of the two counters of $\mathcal{M}$, the locations of the timed automaton to encode the lines of $\mathcal{M}$ and present gadgets that allow a strategy to simulate the unique run of $\mathcal{M}$. If location $k-1$ is reached, then the reachability objective is satisfied. Note that as both ℓ_S and the target location $k-1$ are sinks, a strategy satisfying reachability has to avoid reaching ℓ_S.

We begin by introducing an auxiliary gadget (shown in Figure 3), which takes the first action from the schedule (if necessary, after some time has passed in location ℓ) and reinserts it immediately at the end of the schedule. In $\mathfrak{T}(\mathcal{D})$, the edges in the gadget induce a (unique) run from a state (ℓ, v, σ) with a well-timed $\sigma = (a_0, t_0) \cdots (a_{k-1}, t_{k-1})$ with $a_0 = a$: In ℓ, time has to pass for $t_0 > 0$ units of time, which implies that the last time stamp in the resulting schedule is strictly smaller than 1. Then $a_0 = a$ must be processed (which resets clock x). Thus, location m has to be left immediately (due to the invariant $x \leq 0$) via the edge labeled $\mathsf{schd}(a, 1)$ which adds $(a, 1)$ to the schedule and leads to location ℓ'. Note that the first timestamp in the resulting schedule is $t_1 - t_0$, which is greater than zero due to σ being well-timed. Hence, the resulting schedule σ' (when reaching ℓ') is well-timed as well. Furthermore, σ and σ' contain the same actions with the same multiplicity, i.e., they encode the same counter values. Also, the dashed edges leading to ℓ_S are not enabled when m is reached, as it has to be left immediately, while the next action in the schedule has a non-zero timestamp due to well-timedness. These edges are just drawn for completeness to satisfy deadlock freedom. In the following, we use an edge with label $\mathsf{reschd}(a)$ leading from ℓ to ℓ' to represent the gadget from Figure 3.
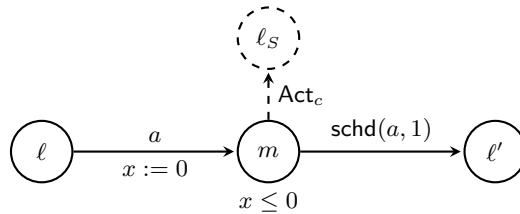

Fig. 3. The gadget $\mathsf{reschd}(a)$ for some $a \in \mathsf{Act}_c$.

Furthermore, in Figure 4, we present the initialization gadget, which adds the termination symbol $\#$ to the (initially) empty schedule on an edge leading from the initial location ℓ_0 to the location 0 representing the initial line of $\mathcal{M}$. Due to the guard $x \leq 0$ on the edge, this action has to be scheduled immediately, i.e., there is no choice to make for a strategy. Thus, there is a unique state reachable from the initial state of $\mathcal{D}_{\mathcal{M}}$ via a single transition, which has a well-timed and well-sorted schedule and encodes the initial configuration of $\mathcal{M}$. Again, the dashed edges to the sink location are just for deadlock-freedom, they are never enabled in the initial state, as the initial schedule is empty.
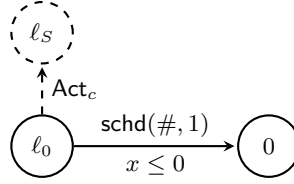

Fig. 4. The initialization gadget.

In the following, we present the gadgets simulating the instructions of a two-counter machine, i.e., increment, decrement, and jump for counter $\mathtt{X}_0$. The gadgets for counter $\mathtt{X}_1$ (and their explanations) are analogous. Afterwards, we present the gadgets for the stop location and the sink ℓ_S.
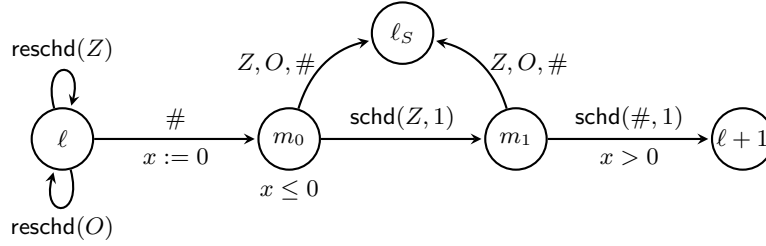

Fig. 5. The gadget for the instruction $\ell : \mathtt{INC}(\mathtt{X}_0)$.

In $\mathfrak{T}(\mathcal{D})$, the edges in the increment-gadget (shown in Figure 5) induce runs of the following form from a state (ℓ, v, σ) with a well-timed *and* well-sorted σ: First, all actions Z and O are rescheduled, resulting in a well-timed schedule σ' with first action $\#$. Then, time passes until the $\#$ in σ' needs to be executed (which removes it from the schedule), leading to location m_0 while resetting x. As a nonzero amount of time passes in ℓ, the last timestamp of the resulting schedule σ'' is strictly smaller than 1. Furthermore, the first action in σ'' has a delay that is still greater than 0, due to well-timedness of σ. Then, due to

the invariant $x \leq 0$ of m_0, a Z has to be added at the end of σ'', leading to location m_1. This implies that the resulting schedule σ''' is well-timed as well, as the timestamp before the newly added Z is strictly smaller than 1, as just argued. Now, due to the guard $x > 0$ on the edge from m_1 to $\ell+1$, some time (determined by the strategy) has to pass before the edge to $\ell+1$ can be taken. This time has to be smaller than the time stamp of the first action in σ''', as this otherwise leads to the sink ℓ_S, implying that the strategy does not satisfy reachability. But if the edge is taken early enough, then the location $\ell + 1$ is reached with a well-timed and well-sorted schedule that has one more Z than σ and the same number of O's as σ, i.e., the increment of counter X_0 is correctly simulated. Note that these runs only differ in the time that is passed in location m_1 and a strategy selects one such run, and can always pick one that leads to $\ell + 1$ and not to the sink. This is indeed the only place where a strategy has a nontrivial choice in our construction.
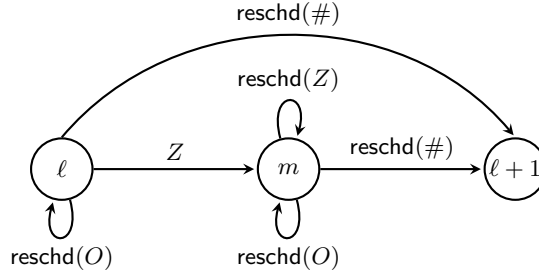

Fig. 6. The gadget for the instruction $\ell : \mathtt{DEC(X_0)}$.

The edges of the decrement-gadget (shown in Figure 6) induce the following (unique) run from a state (ℓ, v, σ) with a well-timed *and* well-sorted σ. If σ contains no Z, i.e., the counter value for X_0 encoded by σ is zero, then all O's in σ are rescheduled (using the self-loop on ℓ) and then the $\#$ is rescheduled (on the edge from ℓ to $\ell+1$). The resulting schedule is well-timed (as the reschedules preserve well-timedness), well-sorted, and has the same number of $O's$ as σ (and still no Z). Thus, the decrement of counter X_0 (holding value zero) is correctly simulated.

On the other hand, if σ contains a Z, i.e., the counter value for X_0 encoded by σ is non-zero, then all O's in σ appearing before the first Z are rescheduled (using the self-loop on ℓ), the first Z is removed from the schedule (using the edge from ℓ to m), thereby ensuring that there is indeed a Z in σ, the remaining Z's and O's are rescheduled (using the self-loop on m), and finally the $\#$ is rescheduled (using the edge from m to $\ell + 1$). As each reschedule and the processing of the first Z preserves well-timedness, the resulting schedule (when reaching $\ell + 1$) is well-timed, and it is well-sorted. Furthermore, it has one less Z than σ and the

same number of O's as σ. Thus, the decrement of counter $\mathtt{X}_0$ (holding a non-zero value) is correctly simulated.
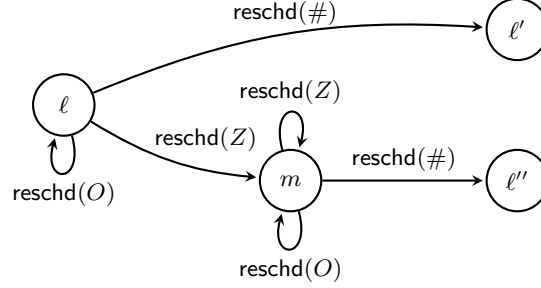
Fig. 7. The gadget for the instruction $\ell : \mathtt{IF}\ \mathtt{X}_0\mathtt{=0}\ \mathtt{GOTO}\ \ell'\ \mathtt{ELSE}\ \mathtt{GOTO}\ \ell''$.

The edges of the jump-gadget (shown in Figure 7) induce the following (unique) run from a state (ℓ, v, σ) with a well-timed *and* well-sorted σ. If σ contains no Z, i.e., the counter value for $\mathtt{X}_0$ encoded by σ is zero, then all O's in σ are rescheduled (using the self-loop on ℓ) and then the $\#$ is rescheduled (on the edge from ℓ to ℓ'). The resulting schedule is well-timed (as the reschedules preserve well-timedness), well-sorted, and has the same number of $O's$ as σ (and still no Z) and thus the jump testing counter $\mathtt{X}_0$ (holding value zero) is correctly simulated.

On the other hand, if σ contains a Z, i.e., the counter value for $\mathtt{X}_0$ encoded by σ is non-zero, then all O's in σ appearing before the first Z are rescheduled (using the self-loop on ℓ), the first Z is rescheduled (using the edge from ℓ to m), thereby ensuring that there is indeed a Z in σ, the remaining Z's and O's are rescheduled (using the self-loop on m), and finally the $\#$ is rescheduled (using the edge from m to ℓ''). As each reschedule preserves well-timedness, the resulting schedule (when reaching ℓ'') is well-timed, and it is well-sorted. Thus, the jump testing counter $\mathtt{X}_0$ (holding a non-zero value) is correctly simulated.

Fig. 8. The gadget for the instruction $k - 1 : \mathtt{STOP}$ (left) and the sink state (right).

Finally, the location $k-1$, corresponding to the line of the stop-instruction, is a sink equipped with a self-loop, as is the sink location (both shown in Figure 8).

Thus, there is always an infinite run starting in these locations, independently of the clock valuation and the schedule.

Now, the timed game under delayed control $\mathcal{D}_{\mathcal{M}}$ consists of the gadgets described above for each line of $\mathcal{M}$ as well as the initialization gadget and the sink gadget. Applying the arguments laid out above inductively shows that a strategy in $\mathcal{D}_{\mathcal{M}}$ simulates the run of $\mathcal{M}$ and that there is a strategy satisfying reachability (of location $k - 1$) if and only if $\mathcal{M}$ terminates.   □

Note that the automaton underlying $\mathcal{D}_{\mathcal{M}}$ only uses a single clock (which is only used locally to avoid scheduling two actions at the same time), has no uncontrollable actions, and a single delay in T. Furthermore, the only decision a strategy has to make is to determine how much time passes in location m_1 in the simulation of an increment.

Furthermore, the reduction can easily be adapted to show that the existence of a strategy satisfying safety is also undecidable. To this end, one just needs to put all locations but ℓ_S and $k - 1$ into F (thereby capturing nontermination of $\mathcal{M}$, which is also undecidable). Finally, the existence of a strategy satisfying both safety and time-divergence is also undecidable, as the additional requirement of time-divergence can be achieved by the strategy: if the infinite run is simulated, then time diverges as simulating each instruction takes one unit of time.

5 Decidability for Bounded Schedules

As timed games under delayed control with unbounded schedules are undecidable, we are now studying timed games under delayed control with bounded schedules. Before we do so, we need to adapt the semantics of timed games under delayed control to disable scheduling actions when the schedule is already full. Say we bound the size of the schedule by cap. Then, a new action can only be scheduled when the schedule contains less than cap scheduled actions. Formally, we redefine the scheduling of (control) actions in the definition of $\mathfrak{T}(\mathcal{D})$ for a timed game under delayed control $\mathcal{D}$ as follows (while leaving the three other cases of the definition unchanged):

– For $a \in \mathsf{Act}_c$ and $t \in T$, there is a transition

$$(\ell, v, (a_0, t_0) \cdots (a_{k-1}, t_{k-1})) \xrightarrow{\mathsf{schd}(a,t)}$$
$$(\ell', v', (a_0, t_0) \cdots (a_{i-1}, t_{i-1})(a, t)(a_i, t_i) \cdots (a_{k-1}, t_{k-1}))$$

if $k < cap$, $t_{i-1} \leq t \leq t_i$, and there is an edge $(\ell, g, \mathsf{schd}(a,t), X', \ell') \in E$ such that $v \models g$, $v' = v[X']$, and $v' \models \mathsf{Inv}(\ell')$.

The definitions of strategies, their outcomes, and them satisfying winning conditions are also unchanged.

Let $\mathcal{D} = (\mathcal{A}, \mathsf{Act}_c, \mathsf{Act}_u, T, F)$ be a timed game under delayed control with $\mathcal{A} = (L, \ell_0, \mathsf{Act}, X, E, \mathsf{Inv})$, where $\mathsf{Act} = \mathsf{Act}_c \cup \mathsf{Act}_u \cup \mathsf{Schd}(\mathsf{Act}_c, T)$. Further, let $cap \in \mathbb{N}$ be the bound on the size of schedules. We define $C = \{0, 1, \ldots, cap - 1\}$.

Intuitively, we will reduce timed games under delayed control with bounded schedules to timed games by storing the schedule using the locations *and* additional clocks. As the size of schedules is bounded by *cap*, this is possible with *cap* many fresh clocks and an exponential blowup of the number of locations. Intuitively, a schedule $\sigma = (a_0, t_0) \cdots (a_{k-1}, t_{k-1})$ with $k \leq cap$ is encoded by three partial functions and a collection $\{x_i \mid i \in C\}$ of fresh clocks as follows:

- The function cl injectively maps indices $j \in \{0, 1, \ldots, k-1\}$ to a clock in $\{x_i \mid i \in C\}$ keeping track of the time since the j-th entry of the schedule has been scheduled. Note that this is different from the value t_j, which keeps track of the time that passes before a_j has to be executed.
- The function ac maps indices $j \in \{0, 1, \ldots, k-1\}$ to the control action a_j of the schedule σ.
- The function dl maps indices $j \in \{0, 1, \ldots, k-1\}$ to elements in T, intended to store the delay the action a_j was originally scheduled with. Then, one can obtain t_j as the difference between $\mathsf{dl}(j)$ and $v(\mathsf{cl}(j))$, provided $\mathsf{cl}(j)$ was last reset when a_j was scheduled.

All three functions are undefined for $j > k - 1$, i.e., they have the same domain $\{0, 1, \ldots, k-1\}$.

We now formalize this intuition. To this end, we write $\mathsf{dom}(f)$ for the domain of a partial function f and $\mathsf{ran}(f)$ for the range $\{f(j) \mid j \in \mathsf{dom}(f)\}$ of f. We say that three partial functions $\mathsf{cl} \colon C \to \{x_i \mid i \in C\}$, $\mathsf{ac} \colon C \to \mathsf{Act}_c$, and $\mathsf{dl} \colon C \to T$ are compatible if $\mathsf{dom}(\mathsf{cl}) = \mathsf{dom}(\mathsf{ac}) = \mathsf{dom}(\mathsf{dl}) = \{0, 1, \ldots, k-1\}$ for some $k \leq cap$ (note that the domains are empty for $k = 0$). Let

$$S = \{(\mathsf{cl}, \mathsf{ac}, \mathsf{dl}) \mid \mathsf{cl}, \mathsf{ac}, \mathsf{dl} \text{ are compatible and } \mathsf{cl} \text{ is injective}\},$$

which is finite.

We define the (delay-free) timed game $\mathcal{G} = (\mathcal{A}^f, \mathsf{Act}_u^f, \mathsf{Act}_c^f, F^f)$ with $\mathcal{A}^f = (L^f, \ell_0^f, \mathsf{Act}^f, X^f, E^f, \mathsf{Inv}^f)$ where

- $L^f = L \times S$,
- $\ell_0^f = (\ell_0, \emptyset, \emptyset, \emptyset)$, where $\emptyset$ is the function with empty domain,
- $\mathsf{Act}^f = \mathsf{Act}$ with $\mathsf{Act}_u^f = \mathsf{Act}_u \cup \mathsf{Act}_c$ and $\mathsf{Act}_c^f = \mathsf{Schd}(\mathsf{Act}_c, T)$ (actions in Act_c are uncontrollable here, as only the scheduling of actions is controllable in a timed game under delayed control),
- $X^f = X \cup \{x_i \mid i \in C\}$,
- $\mathsf{Inv}^f(\ell, \mathsf{cl}, \mathsf{ac}, \mathsf{dl}) = \mathsf{Inv}(\ell) \wedge \bigwedge_{i \in \mathsf{dom}(\mathsf{cl})} \mathsf{cl}(i) \leq \mathsf{dl}(i)$ (the additional invariants ensure that time can pass in a location only until the first scheduled action is due to be executed),
- $F^f = F \times S$,
- and E^f is defined as follows:
    - Let $(\ell, g, a, X', \ell') \in E$ with $a \in \mathsf{Act}_u$ and let $(\mathsf{cl}, \mathsf{ac}, \mathsf{dl}) \in S$. Then, we have $((\ell, \mathsf{cl}, \mathsf{ac}, \mathsf{dl}), g, a, X', (\ell', \mathsf{cl}, \mathsf{ac}, \mathsf{dl})) \in E^f$: Uncontrollable actions can be executed as in $\mathcal{D}$, and they do not influence the encoded schedule.

- Let $(\ell, g, a, X', \ell') \in E$ with $a \in \mathsf{Act}_c$ and let $(\mathsf{cl}, \mathsf{ac}, \mathsf{dl}) \in S$ with $\mathsf{ac}(0) = a$. Then, we have $((\ell, \mathsf{cl}, \mathsf{ac}, \mathsf{dl}), g \wedge g', a, X', (\ell', \mathsf{cl}', \mathsf{ac}', \mathsf{dl}')) \in E^f$ where g' is the constraint $\mathsf{cl}(0) = \mathsf{dl}(0)$, and $\mathsf{cl}'(i) = \mathsf{cl}(i+1)$, $\mathsf{ac}'(i) = \mathsf{ac}(i+1)$, and $\mathsf{dl}'(i) = \mathsf{dl}(i+1)$ for all $i \in \mathsf{dom}(\mathsf{cl}) \setminus \{\max \mathsf{dom}(\mathsf{cl})\}$ (and $\mathsf{cl}', \mathsf{ac}', \mathsf{dl}'$ are undefined for all other inputs): The first action scheduled in the encoded schedule can be executed, provided the right amount of time has passed (captured by the guard g'). This removes the first item from the encoded schedule, which requires to move every other scheduled action "one place to the left".

- Let $(\ell, g, \mathsf{schd}(a, t), X', \ell') \in E$, $(\mathsf{cl}, \mathsf{ac}, \mathsf{dl}) \in S$ such that $\mathsf{dom}(\mathsf{cl}) = \{0, 1, \ldots, k-1\}$ with $k < cap$, and $i^* \in \{0, 1, \ldots, k\}$. Due to $k < cap$, $x^* = x_{\min\{C \setminus \mathsf{ran}(\mathsf{cl})\}}$ is well-defined. We have $((\ell, \mathsf{cl}, \mathsf{ac}, \mathsf{dl}), g \wedge g', a, X' \cup X'', (\ell', \mathsf{cl}', \mathsf{ac}', \mathsf{dl}')) \in E^f$ where, for $i^* > 0$, g' is the guard

$$(\mathsf{cl}(i^* - 1) \geq t - \mathsf{dl}(i^* - 1)) \wedge (\mathsf{cl}(i^*) \leq t - \mathsf{dl}(i^*)),$$

which is equivalent to

$$\mathsf{dl}(i^* - 1) - \mathsf{cl}(i^* - 1) \leq t \leq \mathsf{dl}(i^*) - \mathsf{cl}(i^*).$$

For $i^* = 0$, g' is the guard $(\mathsf{cl}(i^*) \leq t - \mathsf{dl}(i^*))$, which is equivalent to $t \leq \mathsf{dl}(i^*) - \mathsf{cl}(i^*)$. Furthermore, $X'' = \{x_{i^*}\}$ and

$$\mathsf{cl}'(i) = \begin{cases} \mathsf{cl}(i) & \text{if } i < i^*, \\ x^* & \text{if } i = i^*, \\ \mathsf{cl}(i-1) & \text{if } i > i^*, \end{cases}$$

$$\mathsf{ac}'(i) = \begin{cases} \mathsf{ac}(i) & \text{if } i < i^*, \\ a & \text{if } i = i^*, \\ \mathsf{ac}'(i-1) & \text{if } i > i^*, \end{cases}$$

and

$$\mathsf{dl}'(i) = \begin{cases} \mathsf{dl}(i) & \text{if } i < i^*, \\ t & \text{if } i = i^*, \\ \mathsf{dl}(i-1) & \text{if } i > i^*. \end{cases}$$

Here, a new action is scheduled. This action is then inserted into the encoded schedule at any valid index i^*, which requires to move some scheduled actions "one place to the right". Note that the guard g' ensures that the action is inserted at the "right" indices only. Furthermore, the clock x^* assigned to this action is reset, so that it keeps track of the time passed since the action was scheduled.

Lemma 1. *There is a strategy satisfying reachability in the timed game under delayed control $\mathcal{D}$ with schedules bounded by cap if and only if there is a strategy satisfying reachability in the resulting timed game $\mathcal{G}$.*

Proof sketch. An induction shows that there is a bijection between (finite and infinite, respectively) runs of $\mathcal{D}$ with schedules bounded by *cap* and (finite and infinite, respectively) runs of $\mathcal{G}$. This bijection allows one to translate a strategy for $\mathcal{D}$ into a strategy for $\mathcal{G}$ and vice versa. Note that these translations preserve the satisfaction of reachability. □

Now, we are able to state our main result of this section.

Theorem 2. *The following problem is in 2ExpTime: Given a timed game under delayed control and a bound on the size of schedules (encoded in unary), is there a strategy that satisfies reachability?*

Proof. Given a timed game under delayed control $\mathcal{D}$ and a bound *cap*, the timed game $\mathcal{G}$ has *cap* additional clocks and the number of locations blows up by a factor of

$$|S| \leq (cap + 1)^{cap} \cdot (|\mathsf{Act}_c| + 1)^{cap} \cdot (|T| + 1)^{cap}$$
$$= 2^{cap(\log(cap+1) + \log(|\mathsf{Act}_c|+1) + \log(|T|+1))},$$

which is exponential in the number of locations of $\mathcal{A}$ and in *cap* (note that this relies on *cap* being encoded in unary).

Thus, one can construct $\mathcal{G}$ in exponential time and then determine in exponential time (in the size of $\mathcal{G}$) whether there is a strategy for $\mathcal{G}$ satisfying reachability (see Proposition 1). This solves our original problem in doubly-exponential time, due to Lemma 1. □

Again similar results can be obtained for other winning conditions, e.g., safety (with and without time-divergence), parity, etc., as the set of locations of $\mathcal{G}$ is the product of the set of locations of $\mathcal{D}$ and the set S and Lemma 1 holds for all of these conditions.

Remark 2. There are sound, algorithmic methods that, for a given timed game under delayed control $\mathcal{D}$, may produce a bound B (possibly ∞), such that the existence of a strategy for $\mathcal{D}$ satisfying reachability with schedules bounded by B is equivalent to the existence of a strategy satisfying reachability with unbounded schedules (similar claims hold for other winning conditions, e.g., safety). Clearly, due to the undecidability of the latter problem, the methods for producing such bounds B can be sound, but can never be complete, meaning they cannot always compute a (finite) bound if there is one.

One such method is obtained by examining all (syntactic) simple cycles of $\mathcal{D}$: Assume that the time for completing each such cycle is at least T, e.g., if for each such cycle there is a clock x which is reset on the cycle and where $x \geq T$ appears in a guard of the cycle. Also, assume that on each simple cycle at most S actions are scheduled. Finally, assume that D is the maximum delay that appears in any scheduled action. Then, we can conclude that $B = (\lceil \frac{D}{T} \rceil + 1) \cdot S$ is a valid bound B.

Example 4. Reconsider the timed game under delayed control from Figure 2 constituted by the product of $\mathcal{P}$ and $\mathcal{S}$. The only simple cycles involve one of the looping transitions scheduling the kick action with a delay of either 13 or 22. In both cases the minimum time for completing the cycle is 8. Thus, $(\lceil \frac{22}{8} \rceil + 1) \cdot 1 = 4$ is an upper bound on the number of controllable actions that may simultaneously be scheduled.

6   Timed Games are Timed Games under Delayed Control

In this section, we show that our generalization of timed games to timed games under delayed control is conservative in the sense that each timed game (without delay) can be turned into an equivalent timed game under delayed control.

Let $\mathcal{G} = (\mathcal{A}, \mathsf{Act}_c, \mathsf{Act}_u, F)$ be a timed game with $\mathcal{A} = (L, \ell_0, \mathsf{Act}, X, E, \mathsf{Inv})$, i.e., $\mathsf{Act} = \mathsf{Act}_c \cup \mathsf{Act}_u$ and $F \subseteq L$. We define the timed game under delayed control $\mathcal{D} = (\mathcal{A}^d, \mathsf{Act}_c, \mathsf{Act}_u, T^d, F^d)$ with $\mathcal{A}^d = (L^d, \ell_0^d, \mathsf{Act}^d, X, E^d, \mathsf{Inv}^d)$ where

- $L^d = L \times \{0, 1\}$,
- $\ell_0^d = (\ell_0, 0)$,
- $\mathsf{Act}^d = \mathsf{Act}_c \cup \mathsf{Act}_u \cup \mathsf{Schd}(\mathsf{Act}_c, T^d)$,
- $\mathsf{Inv}^d(\ell, 0) = \mathsf{Inv}(\ell)$ and $\mathsf{Inv}^d(\ell, 1) = \mathtt{true}$,
- $T^d = \{0\}$,
- $F^d = F \times \{0, 1\}$, and
- E^d contains the following edges:
    - Let $(\ell, g, a, X', \ell') \in E$ with $a \in \mathsf{Act}_u$. Then, $((\ell, 0), g, a, X', (\ell', 0)) \in E^d$. Intuitively, an uncontrolled edge is simulated using locations of the form $L \times \{0\}$.
    - Let $(\ell, g, a, X', \ell') \in E$ with $a \in \mathsf{Act}_c$. In this case, we have the two edges $((\ell, 0), g, \mathsf{schd}(a, 0), X', (\ell, 1))$ and $((\ell, 1), \mathtt{true}, a, \emptyset, (\ell', 0))$ in E^d. Intuitively, a controlled edge is simulated by checking the guard and applying the reset, but instead of moving to the new location ℓ', the control action is scheduled (with delay 0) and the auxiliary location $(\ell, 1)$ is reached. From here, the only edges available are those that execute the scheduled action, leading to the location $(\ell', 0)$. As no time may pass in $(\ell, 1)$, as the action was scheduled with delay 0, the guard does not have to be checked again and the clocks in X' do not have to be reset again.

Formally, this automaton does *not* satisfy our deadlock-freedom requirement for timed games under delayed control, but adding a sink as described in Remark 1 fixes that issue. Note that the new edges ensuring deadlock-freedom will never be enabled: The size of schedules in $\mathcal{D}$ is always bounded by 1, as locations of the form $(\ell, 1)$, which are reached when scheduling an action, do not allow to schedule further actions and can only be left when the scheduled action is executed.

Theorem 3. *Every timed game $\mathcal{G}$ can be transformed (in linear time) into a timed game under delayed control $\mathcal{D}$ that is equivalent in the following sense: There is a strategy satisfying reachability in $\mathcal{G}$ if and only if there is a strategy satisfying reachability in $\mathcal{D}$.*

Proof sketch. An induction shows that there is a bijection between (finite and infinite, respectively) runs of $\mathcal{G}$ and (finite and infinite, respectively) runs of $\mathcal{D}$. This bijection allows one to translate a strategy for $\mathcal{G}$ into a strategy for $\mathcal{D}$ and vice versa. Note that these translations preserve the satisfaction of reachability.
□

The following corollary follows now immediately from the ExpTime-hardness of determining the existence of strategies satisfying reachability in timed games (see Proposition 1).

Corollary 1. *The following problem is ExpTime-hard: Given a timed game under delayed control whose schedules are of size at most 1, is there a strategy that satisfies reachability?*

Again similar results can be obtained for other winning conditions, e.g., safety (with and without time-divergence), parity, etc., as the set of locations of $\mathcal{D}$ is the product of the set of locations of $\mathcal{G}$ and the set $\{0, 1\}$ and Lemma 3 holds for all of these conditions.

7   Conclusion

Following the persistent and influential work of Martin Fränzle capturing and handling delay inherent to cyber-physical systems in the formal verification of such systems, we introduced timed games where controllable actions do not take effect immediately, but only after some delay. This generalizes (discrete) games under delayed control, as introduced by Martin Fränzle and his collaborators, to the setting of real-time.

In general, solving such games is undecidable, as the number of scheduled actions can be unbounded. On the other hand, we show that this is the only reason for undecidability: timed games under delayed control with bounded schedules can be solved in doubly-exponential time by a reduction to delay-free timed games. Finally, we show that timed games under delayed control are a conservative extension of timed games, i.e., every timed game can be turned into an equivalent timed game under delayed control. As the blow-up of this construction is polynomial, the ExpTime-hardness of solving timed games directly transfers to the problem of solving timed games under delayed control. Note that this leaves a complexity gap: solving timed games under delayed control is ExpTime-hard and in 2ExpTime.

References

1. Alur, R., Dill, D.L.: A theory of timed automata. Theor. Comput. Sci. **126**(2), 183–235 (1994). https://doi.org/10.1016/0304-3975(94)90010-8

2. Asarin, E., Maler, O., Pnueli, A.: Symbolic controller synthesis for discrete and timed systems. In: Antsaklis, P.J., Kohn, W., Nerode, A., Sastry, S. (eds.) Hybrid Systems 1994. LNCS, vol. 999, pp. 1–20. Springer (1994). https://doi.org/10.1007/3-540-60472-3_1

3. Behrmann, G., Cougnard, A., David, A., Fleury, E., Larsen, K.G., Lime, D.: UPPAAL-Tiga: Time for playing games! In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 121–125. Springer (2007). https://doi.org/10.1007/978-3-540-73368-3_14

4. Bengtsson, J., Larsen, K.G., Larsson, F., Pettersson, P., Yi, W.: UPPAAL in 1995. In: Margaria, T., Steffen, B. (eds.) TACAS 1996. LNCS, vol. 1055, pp. 431–434. Springer (1996). https://doi.org/10.1007/3-540-61042-1_66

5. Cassez, F., David, A., Fleury, E., Larsen, K.G., Lime, D.: Efficient on-the-fly algorithms for the analysis of timed games. In: Abadi, M., de Alfaro, L. (eds.) CONCUR 2005. LNCS, Springer (2005). https://doi.org/10.1007/11539452_9

6. Cassez, F., David, A., Larsen, K.G., Lime, D., Raskin, J.: Timed control with observation based and stuttering invariant strategies. In: Namjoshi, K.S., Yoneda, T., Higashino, T., Okamura, Y. (eds.) ATVA 2007. LNCS, vol. 4762, pp. 192–206. Springer (2007). https://doi.org/10.1007/978-3-540-75596-8_15

7. Chaochen, Z., Hoare, C.A.R., Ravn, A.P.: A calculus of durations. Inf. Process. Lett. **40**(5), 269–276 (1991). https://doi.org/10.1016/0020-0190(91)90122-X

8. Chen, M., Fränzle, M., Li, Y., Mosaad, P.N., Zhan, N.: Indecision and delays are the parents of failure - taming them algorithmically by synthesizing delay-resilient control. Acta Informatica **58**(5), 497–528 (2021). https://doi.org/10.1007/S00236-020-00374-7

9. Córdoba, F.C., Palmisano, A., Fränzle, M., Bloem, R., Könighofer, B.: Safety shielding under delayed observation. In: Koenig, S., Stern, R., Vallati, M. (eds.) ICAPS 2023. pp. 80–85. AAAI Press (2023). https://doi.org/10.1609/ICAPS.V33I1.27181

10. Fränzle, M.: Synthesizing controllers from duration calculus. In: Jonsson, B., Parrow, J. (eds.) FTRTFT 1996. LNCS, vol. 1135, pp. 168–187. Springer (1996). https://doi.org/10.1007/3-540-61648-9_40

11. Fränzle, M.: Model-checking dense-time duration calculus. Formal Aspects Comput. **16**(2), 121–139 (2004). https://doi.org/10.1007/S00165-004-0032-Y

12. Fränzle, M., Grosen, T.M., Larsen, K.G., Zimmermann, M.: Monitoring real-time systems under parametric delay. In: Kosmatov, N., Kovács, L. (eds.) IFM 2024. pp. 194–213. LNCS, Springer (2025). https://doi.org/10.1007/978-3-031-76554-4_11

13. Fränzle, M., Kröger, P., Winter, S., Zimmermann, M.: On the existence of reactive strategies resilient to delay. arXiv **2305.19985** (2023). https://doi.org/10.48550/ARXIV.2305.19985, https://doi.org/10.48550/arXiv.2305.19985

14. Fränzle, M., Winter, S., Zimmermann, M.: Strategies resilient to delay: Games under delayed control vs. delay games. In: Achilleos, A., Monica, D.D. (eds.) GandALF 2023. EPTCS, vol. 390, pp. 220–235 (2023). https://doi.org/10.4204/EPTCS.390.14

15. Hosch, F.A., Landweber, L.H.: Finite delay solutions for sequential conditions. In: Nivat, M. (ed.) ICALP 1972. pp. 45–60. North-Holland, Amsterdam (1972)

16. Jurdzinski, M., Trivedi, A.: Reachability-time games on timed automata. In: Arge, L., Cachin, C., Jurdzinski, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596, pp. 838–849. Springer (2007). https://doi.org/10.1007/978-3-540-73420-8_72
17. Klein, F., Zimmermann, M.: How much lookahead is needed to win infinite games? Log. Methods Comput. Sci. **12**(3) (2016). https://doi.org/10.2168/LMCS-12(3:4)2016
18. Klein, F., Zimmermann, M.: Prompt delay. In: Lal, A., Akshay, S., Saurabh, S., Sen, S. (eds.) FSTTCS 2016. LIPIcs, vol. 65, pp. 43:1–43:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016). https://doi.org/10.4230/LIPICS.FSTTCS.2016.43
19. Larsen, K.G., Pettersson, P., Yi, W.: UPPAAL in a nutshell. Int. J. Softw. Tools Technol. Transf. **1**(1-2), 134–152 (1997). https://doi.org/10.1007/S100090050010, https://doi.org/10.1007/s100090050010
20. Maler, O., Pnueli, A., Sifakis, J.: On the synthesis of discrete controllers for timed systems (an extended abstract). In: Mayr, E.W., Puech, C. (eds.) STACS 1995. LNCS, vol. 900, pp. 229–242. Springer (1995). https://doi.org/10.1007/3-540-59042-0_76
21. Minsky, M.L.: Computation: Finite and Infinite Machines. Prentice-Hall (1967)
22. Mosaad, P.N., Fränzle, M., Xue, B.: Temporal logic verification for delay differential equations. In: Sampaio, A., Wang, F. (eds.) ICTAC 2016. LNCS, vol. 9965, pp. 405–421 (2016). https://doi.org/10.1007/978-3-319-46750-4_23
23. Mosaad, P.N., Fränzle, M., Xue, B.: Model checking delay differential equations against metric interval temporal logic. Sci. Ann. Comput. Sci. **27**(1), 77–109 (2017)
24. Tripakis, S.: Decentralized control of discrete-event systems with bounded or unbounded delay communication. IEEE Trans. Autom. Control. **49**(9), 1489–1501 (2004). https://doi.org/10.1109/TAC.2004.834116
25. Zou, L., Fränzle, M., Zhan, N., Mosaad, P.N.: Automatic verification of stability and safety for delay differential equations. In: Kroening, D., Pasareanu, C.S. (eds.) CAV 2015, Part II. LNCS, vol. 9207, pp. 338–355. Springer (2015). https://doi.org/10.1007/978-3-319-21668-3_20