Real-Time Strategy Synthesis for Timed-Arc Petri Net Games via Discretization

Peter Gjøl Jensen, Kim Guldstrand Larsen, and Jiří Srba

Department of Computer Science, Aalborg University,
Selma Lagerlöfs Vej 300, 9220 Aalborg East, Denmark

Abstract. Automatic strategy synthesis for a given control objective can be used to generate correct-by-construction controllers of reactive systems. The existing symbolic approach for continuous timed games is a computationally hard task and current tools like UPPAAL TiGa often scale poorly with the model complexity. We suggest an explicit approach for strategy synthesis in the discrete-time setting and show that even for systems with closed guards, the existence of a safety discrete-time strategy does not imply the existence of a safety continuous-time strategy and vice versa. Nevertheless, we prove that the answers to the existence of discrete-time and continuous-time safety strategies coincide on a practically motivated subclass of urgent controllers that either react immediately after receiving an environmental input or wait with the decision until a next event is triggered by the environment. We then develop an on-the-fly synthesis algorithm for discrete timed-arc Petri net games. The algorithm is implemented in our tool TAPAAL and based on the experimental evidence, we discuss the advantages of our approach compared to the symbolic continuous-time techniques.

1 Introduction

Formal methods and model checking techniques have been traditionally used to verify whether a given system model complies with its specification. However, when we consider formal (game) models where both the controller and the environment can make choices, the question now changes to finding a controller strategy such that any behaviour under such a fixed strategy complies with the given specification. The model checking approach can be used as a try-and-fail technique to check whether a given controller is correct but automatic synthesis of a controller correct-by-construction, as already proposed by Church [13, 12], is a more difficult problem as illustrated by the SYNTCOMP competition and SYNT workshop [1]. This area has recently seen renewed interest, partly given the rise in computational power that makes the synthesis feasible. We focus on the family of timed systems, where for the model of timed automata [2] synthesis has already been proposed [33] and implemented [4, 11].

In the area of model checking, symbolic continuous-time on-the-fly methods were ensuring the success of tools such as Kronos [9], UPPAAL [5], Tina [6]

and Romeo [21], utilizing the zone abstraction approach [2] via the data structure DBM [16]. These symbolic techniques were recently employed in on-the-fly algorithms [28] for synthesis of controllers for timed games [4, 11, 33]. While these methods scale well for classical reachability, the limitation of symbolic techniques is more apparent when used for liveness properties and for solving timed games. We have shown that for reachability and liveness properties, the discrete-time methods performing point-wise exploration of the state-space can prove competitive on a wide range of problems [3], in particular in combination with additional techniques as time-darts [25], constant-reducing approximation techniques [7] and memory-preserving data structures as PTrie [24].

In this paper, we benefit from the recent advances in the discrete-time verification of timed systems and suggest an on-the-fly point-wise algorithm for the synthesis of timed controllers relative to safety objectives (avoiding undesirable behaviour). The algorithm is described for a novel game extension of the well-studied timed-arc Petri net formalism [8, 23] and we show that in the general setting the existence of a controller for a safety objective in the discrete-time setting does not imply the existence of such a controller in the continuous-time setting and vice versa, not even for systems with closed guards—contrary to the fact that continuous-time and discrete-time reachability problems coincide for timed models [10], in particular also for timed-arc Petri nets [30]. However, if we restrict ourselves to the practically relevant subclass of urgent controllers that either react immediately to the environmental events or simply wait for another occurrence of such an event, then we can use the discrete-time methods for checking the existence of a continuous-time safety controller on closed timed-arc Petri nets. The algorithm for controller synthesis is implemented in the tool TAPAAL [15], including the memory optimization technique via PTrie [24], and the experimental data show a promising performance on a large data-set of infinite job scheduling problems as well as on other examples.

Related Work. An on-the-fly algorithm for synthesizing continuous-time controllers for both safety, reachability and time-optimal reachability for time automata was proposed by Cassez et al. [11] and later implemented in the tool UPPAAL TiGa [4]. This work is based on the symbolic verification techniques invented by Alur and Dill [2] in combination with ideas on synthesis by Pnueli et. al [33] and on-the-fly dependency graph algorithms suggested by Liu and Smolka [28]. For timed games, abstraction refinement approaches have been proposed and implemented by Peter et. al [31, 32] and Finkbeiner et. al [19] as an attempt to speed up synthesis, while using the same underlying symbolic representation as UPPAAL TiGa. These abstraction refinement methods are complementary to the work presented here. Our work uses the formalism of timed-arc Petri nets that has not been studied in this context before and we rely on the methods with discrete interpretation of time as presented by Andersen et. al [3]. As an additional contribution, we implement our solution in the tool TAPAAL, utilizing memory reduction techniques by Jensen et. al [24], and compare the performance of both discrete-time and continuous-time techniques.

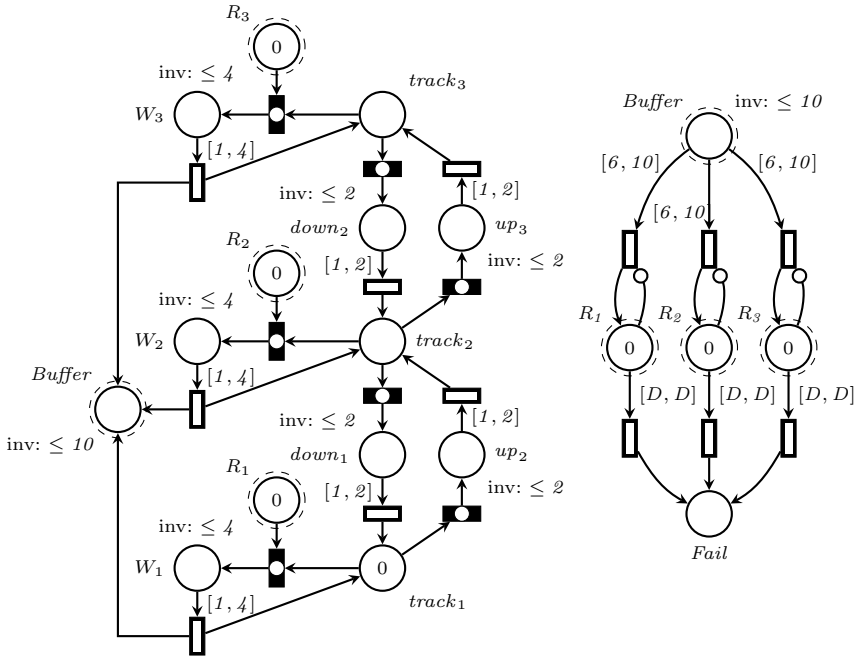

Fig. 1: A timed-arc Petri net game model of a harddisk

Control synthesis and supervisory control was also studied for the family of Petri net models [17, 18, 34, 36] but these works do not consider the timing aspects.

2 Motivating Example of Disk Operation Scheduling

We shall now provide an intuitive description of the timed-arc Petri net game of *disk operation scheduling* in Figure 1, modelling the scheduler of a mechanical harddisk drive (left) and a number of read stream requests (right) that should be fulfilled within a given deadline D. The net consists of *places* drawn as circles (the dashed circle around the places R_1, R_2, R_3 and *Buffer* simply means that these places are shared between the two subnets) and *transitions* drawn as rectangles that are either filled (controllable transitions) or framed only (environmental transitions). Places can contain *tokens* (like the places R_1 to R_3 and the place $track_1$) and each token carries its own age. Initially all token ages are 0. The net also contains *arcs* from places to transitions (input arcs) or transitions to places (output arcs). The input arcs are further decorated with *time intervals* restricting the ages of tokens that can be consumed along the arc. If the time interval is missing, we assume the default $[0, \infty]$ interval not restricting the ages of tokens in any way.

In the initial *marking* (token configuration) depicted in our example, the two transitions connected by input arcs to the place $track_1$ are *enabled* and the controller can decide to *fire* either of them. As the transitions contain a white circle, they are *urgent*, meaning that time cannot pass as long at least one urgent transition is enabled. Suppose now that the controller decides to fire the transition on the left of the place $track_1$. As a result of firing the transition, the two tokens in R_1 and $track_1$ will be consumed and a new token of age 0 produced to the place W_1. Tokens can be also transported via a pair of an input and output *transport arcs* (not depicted in our example) that will transport the token from the input to the output place while preserving its age.

In the new marking we just achieved, no transition is enabled due to the time interval $[1, 4]$ on the input arc of the environmental transition connected to the place W_1. However, after one time unit passes and the token in W_1 becomes of age 1, the transition becomes enabled and the environment may decide to fire it. On the other hand, the place W_1 also contains an *age invariant* ≤ 4, requiring that the age of any token in that place may not exceed 4. Hence after age of the token reaches 4, time cannot progress anymore and the environment is forced to fire the transition, producing two fresh tokens into the places *Buffer* and $track_1$. Hence, reading the data from track 1 of the disk takes between 1ms to 4ms (depending on the actual rotation of the disk) and it is the environment that decides the actual duration of the reading operation.

The idea is that the disk has three tracks (positions of the reading head) and at each track $track_i$ the controller has the choice of either reading the data from the given track (assuming there is a reading request represented by a token in the place R_i) or move the head to one of the neighbouring tracks (such a mechanical move takes between 1ms to 2ms). The reading requests are produced by the subnet on the right where the environment decides when to generate a reading request in the interval between 6ms to 10ms. The number of tokens in the right subnet represents the parallel reading streams. The net also contains *inhibitor arcs* with a cirle-headed tip that prohibit the environmental transitions from generating a reading request on a given track if there is already one. Finally, if the reading request takes too long and the age of the token in R_i reaches the age D, the environment has the option to place a token in the place *Fail*.

The control synthesis problem asks to find a strategy for firing the controllable transitions that guarantees no failure, meaning that irrelevant of the behaviour of the environment, the place *Fail* never becomes marked (safety control objective). The existence of such a control strategy depends on the chosen value of D and the complexity of the controller synthesis problem can be scaled by adding further tracks (in the subnet of the left) or allowing for more parallel reading streams (in the subnet on the right). In what follows, we shall describe how to automatically decide in the discrete-time setting (where time can be increased only by nonnegative integer values) whether a controller strategy exists. As the controllable transitions are urgent in our example, the existence of such a discrete-time control strategy implies also the existence of a continuous-

time control strategy where the environment is free to fire transitions after an arbitrary delay taken from the dense time domain.

3 Definitions

Let $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ and $\mathbb{N}_0^\infty = \mathbb{N}_0 \cup \{\infty\}$. Let $\mathbb{R}^{\geq 0}$ be the set of all nonnegative real numbers. A *timed transition system* (TTS) is a triple $(S, Act, \rightarrow)$ where S is the set of states, Act is the set of actions and $\rightarrow \subseteq S \times (Act \cup \mathbb{R}^{\geq 0}) \times S$ is the transition relation written as $s \xrightarrow{a} s'$ whenever $(s, a, s') \in \rightarrow$. If $a \in Act$ then we call it a *switch transition*, if $a \in \mathbb{R}^{\geq 0}$ we call it a *delay transition*. We also define the set of *well-formed closed time intervals* as $\mathcal{I} \stackrel{\text{def}}{=} \{[a, b] \mid a \in \mathbb{N}_0, b \in \mathbb{N}_0^\infty, a \leq b\}$ and its subset $\mathcal{I}^{\text{inv}} \stackrel{\text{def}}{=} \{[0, b] \mid b \in \mathbb{N}_0^\infty\}$ used in age invariants.

Definition 1 (Timed-Arc Petri Net). *A* timed-arc Petri net *(TAPN) is a 9-tuple* $N = (P, T, T_{urg}, IA, OA, g, w, Type, I)$ *where*

- P *is a finite set of* places,
- T *is a finite set of* transitions *such that* $P \cap T = \emptyset$,
- $T_{urg} \subseteq T$ *is the set of* urgent transitions,
- $IA \subseteq P \times T$ *is a finite set of* input arcs,
- $OA \subseteq T \times P$ *is a finite set of* output arcs,
- $g : IA \rightarrow \mathcal{I}$ *is a* time constraint function *assigning guards to input arcs such that*
  - *if* $(p, t) \in IA$ *and* $t \in T_{urg}$ *then* $g((p, t)) = [0, \infty]$,
- $w : IA \cup OA \rightarrow \mathbb{N}$ *is a function assigning* weights *to input and output arcs,*
- $Type : IA \cup OA \rightarrow \mathbf{Types}$ *is a* type function *assigning a type to all arcs where* $\mathbf{Types} = \{Normal, Inhib\} \cup \{Transport_j \mid j \in \mathbb{N}\}$ *such that*
  - *if* $Type(z) = Inhib$ *then* $z \in IA$ *and* $g(z) = [0, \infty]$,
  - *if* $Type((p, t)) = Transport_j$ *for some* $(p, t) \in IA$ *then there is exactly one* $(t, p') \in OA$ *such that* $Type((t, p')) = Transport_j$,
  - *if* $Type((t, p')) = Transport_j$ *for some* $(t, p') \in OA$ *then there is exactly one* $(p, t) \in IA$ *such that* $Type((p, t)) = Transport_j$,
  - *if* $Type((p, t)) = Transport_j = Type((t, p'))$ *then* $w((p, t)) = w((t, p'))$,
- $I : P \rightarrow \mathcal{I}^{inv}$ *is a function assigning* age invariants *to places.*

Remark 1. Note that for transport arcs we assume that they come in pairs (for each type $Transport_j$) and that their weights match. Also for inhibitor arcs and for input arcs to urgent transitions, we require that the guards are $[0, \infty]$. This restriction is important for some of the results presented in this paper and it also guarantees that we can use DBM-based algorithms in the tool TAPAAL [15].

Before we give the formal semantics of the model, let us fix some notation. Let $N = (P, T, T_{urg}, IA, OA, g, w, Type, I)$ be a TAPN. We denote by ${}^\bullet x \stackrel{\text{def}}{=} \{y \in P \cup T \mid (y, x) \in IA \cup OA,\ Type((y, x)) \neq Inhib\}$ the preset of a transition or a place x. Similarly, the postset is defined as $x^\bullet \stackrel{\text{def}}{=} \{y \in P \cup T \mid (x, y) \in (IA \cup OA)\}$.

Let $\mathcal{B}(\mathbb{R}^{\geq 0})$ be the set of all finite multisets over $\mathbb{R}^{\geq 0}$. A *marking* M on N is a function $M : P \longrightarrow \mathcal{B}(\mathbb{R}^{\geq 0})$ where for every place $p \in P$ and every token $x \in M(p)$ we have $x \in I(p)$, in other words all tokens have to satisfy the age invariants. The set of all markings in a net N is denoted by $\mathcal{M}(N)$.

We write (p, x) to denote a token at a place p with the age $x \in \mathbb{R}^{\geq 0}$. Then $M = \{(p_1, x_1), (p_2, x_2), \ldots, (p_n, x_n)\}$ is a multiset representing a marking M with n tokens of ages x_i in places p_i. We define the size of a marking as $|M| = \sum_{p \in P} |M(p)|$ where $|M(p)|$ is the number of tokens located in the place p.

Definition 2 (Enabledness). *Let $N = (P, T, T_{urg}, IA, OA, g, w, Type, I)$ be a TAPN. We say that a transition $t \in T$ is* enabled *in a marking M by the multisets of tokens $In = \{(p, x_p^1), (p, x_p^2), \ldots, (p, x_p^{w((p,t))}) \mid p \in {}^\bullet t\} \subseteq M$ and $Out = \{(p', x_{p'}^1), (p', x_{p'}^2), \ldots, (p', x_{p'}^{w((t,p'))}) \mid p' \in t^\bullet\}$ if*

- *for all input arcs except the inhibitor arcs, the tokens from In satisfy the age guards of the arcs, i.e.*

$$\forall p \in {}^\bullet t.\ x_p^i \in g((p,t)) \text{ for } 1 \leq i \leq w((p,t))$$

- *for any inhibitor arc pointing from a place p to the transition t, the number of tokens in p is smaller than the weight of the arc, i.e.*

$$\forall (p,t) \in IA.\, Type((p,t)) = Inhib \Rightarrow |M(p)| < w((p,t))$$

- *for all input arcs and output arcs which constitute a transport arc, the age of the input token must be equal to the age of the output token and satisfy the invariant of the output place, i.e.*

$$\begin{aligned} \forall (p,t) \in IA. \forall (t,p') \in OA.\, Type((p,t)) = Type((t,p')) = Transport_j \\ \Rightarrow \left(x_p^i = x_{p'}^i \wedge x_{p'}^i \in I(p')\right) \text{ for } 1 \leq i \leq w((p,t)) \end{aligned}$$

- *for all normal output arcs, the age of the output token is 0, i.e.*

$$\forall (t,p') \in OA.\, Type((t,p')) = Normal \Rightarrow x_{p'}^i = 0 \text{ for } 1 \leq i \leq w((t,p')).$$

A given TAPN N defines a TTS $T(N) \stackrel{\text{def}}{=} (\mathcal{M}(N), T, \rightarrow)$ where states are the markings and the transitions are as follows.

- If $t \in T$ is enabled in a marking M by the multisets of tokens *In* and *Out* then t can *fire* and produce the marking $M' = (M \setminus In) \uplus Out$ where $\uplus$ is the multiset sum operator and $\setminus$ is the multiset difference operator; we write $M \xrightarrow{t} M'$ for this switch transition.
- A time *delay* $d \in \mathbb{R}^{\geq 0}$ is allowed in M if
  - $(x + d) \in I(p)$ for all $p \in P$ and all $x \in M(p)$, and
  - if $M \xrightarrow{t} M'$ for some $t \in T_{urg}$ then $d = 0$.

  By delaying d time units in M we reach the marking M' defined as $M'(p) = \{x + d \mid x \in M(p)\}$ for all $p \in P$; we write $M \xrightarrow{d} M'$ for this delay transition.

Let $\rightarrow \stackrel{\text{def}}{=} \bigcup_{t \in T} \xrightarrow{t} \cup \bigcup_{d \in \mathbb{R}^{\geq 0}} \xrightarrow{d}$. By $M \xrightarrow{d,t} M'$ we denote that there is a marking M'' such that $M \xrightarrow{d} M'' \xrightarrow{t} M'$.

The semantics defined above in terms of timed transition systems is called the *continuous-time semantics*. If we restrict the possible delay transitions to take values only from nonnegative integers and the markings to be of the form $M : P \longrightarrow \mathcal{B}(\mathbb{N}_0)$, we call it the *discrete-time semantics*.

3.1 Timed-Arc Petri Net Game

We shall now extend the TAPN model into the game setting by partitioning the set of transitions into the controllable and uncontrollable ones.

Definition 3 (Timed-Arc Petri Net Game). *A Timed-Arc Petri Net Game (TAPG) is a TAPN with its set of transitions T partitioned into the controller T_{ctrl} and environment T_{env} sets.*

Let G be a fixed TAPG. Recall that $\mathcal{M}(G)$ is the set of all markings over the net G. A *controller strategy* for the game G is a function

$$\sigma : \mathcal{M}(G) \to \mathcal{M}(G) \cup \{wait\}$$

from markings to markings or the special symbol *wait* such that

- if $\sigma(M) = wait$ then either M can delay forever ($M \xrightarrow{d}$ for all $d \in \mathbb{R}^{\geq 0}$), or there is $d \in \mathbb{R}^{\geq 0}$ where $M \xrightarrow{d} M'$ and for all $d'' \in \mathbb{R}^{\geq 0}$ for all $t \in T_{ctrl}$ we have that if $M' \xrightarrow{d''} M''$ then $M'' \stackrel{t}{\not\rightarrow}$, and
- if $\sigma(M) = M'$ then there is $d \in \mathbb{R}^{\geq 0}$ and there is $t \in T_{ctrl}$ where $M \xrightarrow{d,t} M'$.

Intuitively, a controller can in a given marking M either decide to wait indefinitely (assuming that it is not forced by age invariants or urgency to perform some controllable transition) or it can suggest a delay followed by a controllable transition firing. The environment can in the marking M also propose to wait (unless this is not possible due to age invariants or urgency) or suggest a delay followed by firing of an uncontrollable transition. If both the controller and environment propose transition firing, then the one preceding with a shorter delay takes place. In the case where both the controller and the environment propose the same delay followed by a transition firing, then any of these two firings can (nondeterministically) happen. This intuition is formalized in the notion of *plays* following a fixed controller strategy that summarize all possible executions for any possible environment.

Let $\pi = M_1 M_2 \ldots M_n \ldots \in \mathcal{M}(G)^\omega$ be an arbitrary finite or infinite sequence of markings over G and let M be a marking. We define the concatenation of M with π as $M \circ \pi = M M_1 \ldots M_n \ldots$ and extend it to the sets of sequences $\Pi \subseteq \mathcal{M}(G)^\omega$ so that $M \circ \Pi = \{M \circ \pi \mid \pi \in \Pi\}$.

Definition 4 (Plays According to the Strategy σ). *Let G be a TAPG, M a marking on G and σ a controller strategy for G. We define a function $\mathbb{P}_\sigma : \mathcal{M}(G) \to 2^{\mathcal{M}(G)^\omega}$ returning for a given marking M the set of all possible plays starting from M under the strategy σ.*

- *If $\sigma(M) = wait$ then $\mathbb{P}_\sigma(M) = \{M \circ \mathbb{P}_\sigma(M') \mid d \in \mathbb{R}^{\geq 0},\ t \in T_{env},\ M \xrightarrow{d,t} M'\} \cup X$ where $X = \{M\}$ if $M \xrightarrow{d}$ for all $d \in \mathbb{R}^{\geq 0}$, or if there is $d' \in \mathbb{R}^{\geq 0}$ such that $M \xrightarrow{d'} M'$ and $M' \xrightarrow{d''} \not\to$ for any $d'' > 0$ and $M' \xrightarrow{t} \not\to$ for any $t \in T_{env}$, otherwise $X = \emptyset$.*
- *If $\sigma(M) \neq wait$ then according to the definition of controller strategy we have $M \xrightarrow{d,t} \sigma(M)$ and we define $\mathbb{P}_\sigma(M) = \{M \circ \mathbb{P}_\sigma(\sigma(M))\} \cup \{M \circ \mathbb{P}_\sigma(M') \mid d' \leq d, t' \in T_{env}, M \xrightarrow{d',t'} M'\}$.*

The first case says that the plays from the marking M where the controller wants to wait consist either of the marking M followed by any play from a marking M' that can be reached by the environment from M after some delay and firing a transition from T_{env}, or a finite sequence finishing the marking M if it is the case that M can delay forever, or we can reach a deadlock where no further delay is possible and no transition can fire.

The second case where the controller suggests a transition firing after some delay, contains M concatenated with all possible plays from $\sigma(M)$ and from $\sigma(M')$ for any M' that can be reached by the environment before or at the same time the controller suggests to perform its move.

We can now define the safety objectives for TAPGs that are boolean expressions over arithmetic predicates which observe the number of tokens in the different places of the net. Let φ be so a boolean combination of predicates of the form $e \bowtie e$ where $e ::= p \mid n \mid e + e \mid e - e \mid e * e$ and where $p \in P$, $\bowtie \in \{<, \leq, =, \neq, \geq, >\}$ and $n \in \mathbb{N}_0$. The semantics of φ in a marking M is given in the natural way, assuming that p stands for $|M(p)|$ (the number of tokens in the place p). We write $M \models \varphi$ if φ evaluates in the marking M to true. We can now state the safety synthesis problem.

Definition 5 (Safety Synthesis Problem). *Given a marked TAPG G with the initial marking M_0 and a safety objective φ, decide if there is a controller strategy σ such that*

$$\forall \pi \in \mathbb{P}_\sigma(M_0). \forall M \in \pi. M \models \varphi\ . \tag{1}$$

If Equation (1) holds then we say that σ is a winning controller strategy *for the objective φ.*

4 Controller Synthesis in Continuous vs. Discrete Time

It is known that for classical TAPNs the continuous and discrete-time semantics coincide up to reachability [30], which is what safety synthesis reduces to if the

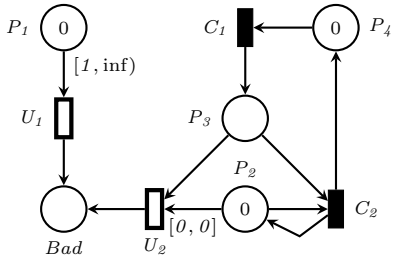

(a) A TAPG where $Bad \leq 0$ can be guaranteed by the controller under the continuous-time semantics but not under the discrete-time semantics by exploiting Zeno behavior.

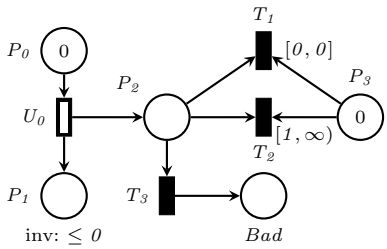

(b) A TAPG where $Bad \leq 0$ can be guaranteed by the controller under the discrete-time semantics but not under the continuous-time semantics.

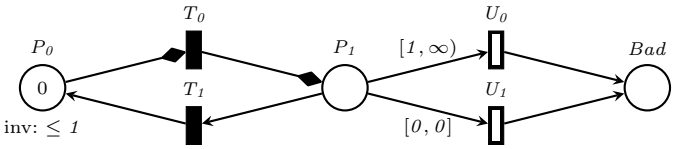

(c) A TAPG where $Bad \leq 0$ can be guaranteed by the controller under the continuous-time semantics but not under the discrete-time semantics (without exploiting Zeno behavior).

Fig. 2: Difference between continuous and discrete-time semantics

set of controllable transitions is empty. Contrary to this, Figures 2a and 2b show that this does not hold in general for safety strategies.

For the game in Figure 2a, there exists a strategy for the controller and the safety objective $Bad \leq 0$ but this is the case only in the continuous-time semantics as the controller has to keep the age of the token in place P_1 strictly below 1, otherwise the environment can mark the place Bad by firing U_1. However, if the controller fires transition C_1 without waiting, U_2 becomes enabled and the environment can again break the safety. Hence it is impossible to find a discrete-time strategy as even the smallest possible delay of 1 time unit will enable U_1. However, if the controller waits an infinitesimal amount (in the continuous semantics) and fires C_1, then U_2 will not be enabled as the token in P_2 aged slightly. The controller can now fire C_2 and repeat this strategy over and over in order to keep the token in P_1 from ever reaching the age of 1.

The counter example described before relies on Zeno behaviour, however, this is not needed if we use transport arcs that do not reset the age of tokens (depicted by arrows with diamond-headed tips), as demonstrated in Figure 2c. Here the only winning strategy for the controller to avoid marking the place Bad is to delay some fraction and then fire T_0. Any possible integer delay (1 or 0) will enable the environment to fire U_0 or U_1 before the controller gets to fire T_1. Hence we get the following lemma.

Lemma 1. *There is a TAPG and a safety objective where the controller has a winning strategy in the continuous-time semantics but not in the discrete-time semantics.*

Figure 2b shows, on the other hand, that a safety strategy guaranteeing $Bad \leq 0$ exists only in the discrete-time semantics but not in the continuous-time one where the environment can mark the place Bad by initially delaying 0.5 and then firing U_0. This will produce a token in P_1 which restricts the time from progressing further and thus forces the controller to fire T_3 as this is the only enabled transition. On the other hand, in the discrete-time semantics the environment can either fire U_0 immediately but then T_1 will be enabled, or it can wait (a minimum of one time unit), however then T_2 will be enabled. Hence the controller can in both cases avoid the firing of T_3 in the discrete-time semantics. This implies the following lemma.

Lemma 2. *There is a TAPG and a safety objective where the controller has a winning strategy in the discrete-time semantics but not in the continuous-time semantics.*

This indeed means that the continuous and discrete-time semantics are incomparable and it makes sense to consider both of them, depending on the concrete application domain and the fact whether we consider discretized or continuous time. Nevertheless, there is a practically relevant subclass of the problem where we consider only urgent controllers and where the two semantics coincide. We say that a given TAPG is with an *urgent controller* if all controllable transitions are urgent, formally $T_{ctrl} \subseteq T_{urg}$.

Theorem 1. *Let G be a TAPG with urgent controller and let φ be a safety objective. There is a winning controller strategy for G and φ in the discrete-time semantics iff there is a winning controller strategy for G and φ in the continuous-time semantics.*

Proof (Sketch). The existence of a winning controller strategy in the continuous-time semantics clearly implies the existence of such a strategy also in the discrete-time because here the environment is restricted to playing only integer delays and the controller can always react to these according to the continuous-time strategy that exists by our assumption. Because the controller is making only urgent choices or waits for the next environmental move, all transitions happen in the discrete-time points.

For the other direction, we prove the converse via the use of linear programming as used e.g. in [30]. Assuming that the urgent controller does not have a winning strategy in the continuous-time semantics, we will argue that the controller does not have a winning strategy in the discrete-time semantics either. Due to the assumption, we know that the environment can in any current marking choose a real-time delay and an uncontrollable transition in such a way that irrelevant of what the controller chooses, it eventually reaches a marking violating the safety condition φ. Such an environmental strategy can be described

as a finite tree where nodes are markings, edges contain the information about the delay and transition firing, the branching describes all controller choices and each leaf of the tree is a marking that satisfies $\neg\varphi$. The existence of this environmental strategy follows from the determinacy of the game that guarantees that one of the players must have a winning strategy (to see this, we realize that the environmental strategy contains only finite branches, all of them ending in a marking satisfying $\neg\varphi$, and hence we have an instance of an open game that is determined by the result of Gale and Stewart [20]—see also [22]).

As we assume that the environment can win in the continuous-time semantics, the delays in the tree may be nonnegative real numbers (controller's moves in the tree are always with delay 0). Our aim is to show that there is another winning tree for the environment, however, with integer delays only. This can be done by replacing the delays in the tree by variables and reformulating the firing conditions of the transitions in the tree as a linear program. Surely, the constraints in the linear program have, by our assumption, a nonnegative real solution. Moreover, the constraint system uses only closed difference constraints (nonstrictly bounding the difference of two variables from below or above) and we can therefore reduce the linear program to a shortest-path problem with integer weights only and this implies that an integer solution exists too [14]. This means that there is a tree describing an environmental winning strategy using only integer delays and hence the controller does not have a winning strategy in the discrete-time setting. The technical details of the proof are provided in the full version of the paper. □

5 Discrete-Time Algorithm for Controller Synthesis

We shall now define the discrete-time algorithm for synthesizing controller strategies for TAPGs. As the state-space of a TAPG is infinite in several aspects (the number of tokens in reachable markings can be unbounded and even for bounded nets the ages of tokens can be arbitrarily large), the question of deciding the existence of a controller strategy is in general undecidable (already the classical reachability is undecidable [35] for TAPNs).

We address the undecidability issue by enforcing a given constant k, bounding the number of tokens in any marking reached by the controller strategy. This means that instead of checking the safety objective φ, we verify instead the safety objective $\varphi_k = \varphi \wedge k \geq \sum_{p \in P} p$ that at the same time ensures that the total number of tokens is at most k. This will, together with the extrapolation technique below, guarantee the termination of the algorithm.

5.1 Extrapolation of TAPGs

We shall now recall a few results from [3] that allow us to make finite abstractions of bounded nets (in the discrete-time semantics). The theorems and lemmas in the rest of this section hold also for continuous-time semantics, however, the finiteness of the extrapolated state space is not guaranteed in this case.

Let $G = (P, T, T_{env}, T_{ctrl}, T_{urg}, IA, OA, g, w, Type, I)$ be a TAPG. In [3] the authors provide an algorithm for computing a function $C_{max} : P \to (\mathbb{N}_0 \cup \{-1\})$ returning for each place $p \in P$ the maximum constant associated to this place, meaning that the ages of tokens in place p that are strictly greater than $C_{max}(p)$ are irrelevant. The function $C_{max}(p)$ for a given place p is computed by essentially taking the maximum constant appearing in any outgoing arc from p and in the place invariant of p, where a special care has to be taken for places with outgoing transport arcs (details are discussed in [3]). In particular, places where $C_{max}(p) = -1$ are the so-called *untimed* places where the age of tokens is not relevant at all, implying that all the intervals on their outgoing arcs are $[0, \infty]$.

Let M be a marking of G. We split it into two markings $M_>$ and $M_\leq$ where $M_>(p) = \{x \in M(p) \mid x > C_{max}(p)\}$ and $M_\leq(p) = \{x \in M(p) \mid x \leq C_{max}(p)\}$ for all places $p \in P$. Clearly, $M = M_> \uplus M_\leq$.

We say that two markings M and M' in the net G are equivalent, written $M \equiv M'$, if $M_\leq = M'_\leq$ and for all $p \in P$ we have $|M_>(p)| = |M'_>(p)|$. In other words M and M' agree on the tokens with ages below the maximum constants and have the same number of tokens above the maximum constant.

The relation $\equiv$ is an equivalence relation and it is also a timed bisimulation (see e.g. [27]) where delays and transition firings on one side can be matched by exactly the same delays and transition firings on the other side and vice versa.

Theorem 2 ([3]). *The relation $\equiv$ is a timed bisimulation.*

We can now define canonical representatives for each equivalence class of $\equiv$.

Definition 6 (Cut). *Let M be a marking. We define its canonical marking $cut(M)$ by $cut(M)(p) = M_\leq(p) \uplus \{ \underbrace{C_{max}(p)+1, \ldots, C_{max}(p)+1}_{|M_>(p)| \text{ times}} \}$.*

Lemma 3 ([3]). *Let M, M_1 and M_2 be markings. Then (i) $M \equiv cut(M)$, and (ii) $M_1 \equiv M_2$ if and only if $cut(M_1) = cut(M_2)$.*

5.2 The Algorithm

After having introduced the extrapolation function *cut* and our enforcement of the k-bound, we can now design an algorithm for computing a controller strategy σ, provided such a strategy exists.

Algorithm 1 describes a discrete-time method to check if there is a controller strategy or not. It is centered around four data structures: *Waiting* for storing markings to be explored, *Losing* that contains marking where such a strategy does not exist, *Depend* for maintaining the set of dependencies to be reinserted to the waiting list whenever a marking is declared as losing, and *Processed* for already processed markings. All markings in the algorithm are always considered modulo the *cut* extrapolation. The algorithm performs a forward search by repeatedly selecting a marking M from *Waiting* and if it can determine that

Algorithm 1: Safety Synthesis Algorithm

Input: A TAPG $G = (P, T, T_{env}, T_{ctrl}, T_{urg}, IA, OA, g, w, Type, I)$, initial marking M_0, a safety objective φ and a bound k.

Output: *tt* if there exists a controller strategy ensuring φ and not exceeding k tokens in any intermediate marking, *ff* otherwise

begin
  $Waiting := Losing := Processed = \emptyset$; $\varphi_k = \varphi \wedge k \geq \sum_{p \in P} p$;
  $M \leftarrow cut(M_0)$; $Depend[M] \leftarrow \emptyset$;
  if $M \not\models \varphi_k$ **then**
    $Losing \leftarrow \{M\}$
  else
    $Waiting \leftarrow \{M\}$
  while $Waiting \neq \emptyset \wedge cut(M_0) \notin Losing$ **do**
    $M \leftarrow pop(Waiting)$;
    $Succs_{env} := \{cut(M') \mid t \in T_{env},\ M \xrightarrow{t} M'\}$;
    $Succs_{ctrl} := \{cut(M') \mid t \in T_{ctrl},\ M \xrightarrow{t} M'\}$;
    $Succs_{delay} := \begin{cases} \emptyset & \text{if } M \not\xrightarrow{1} \\ \{cut(M')\} & \text{if } M \xrightarrow{1} M' \end{cases}$
    if $\exists M' \in Succs_{env}$ *s.t.* $M' \not\models \varphi_k \vee M' \in Losing$ **then**
      $Losing \leftarrow Losing \cup \{M\}$;
      $Waiting \leftarrow (Waiting \cup Depend[M]) \setminus Losing$;
    else
      if $Succs_{ctrl} \cup Succs_{delay} \neq \emptyset \wedge \forall M' \in Succs_{ctrl} \cup Succs_{delay}.$ $M' \not\models \varphi_k \vee M' \in Losing$ **then**
        $Losing \leftarrow Losing \cup \{M\}$;
        $Waiting \leftarrow (Waiting \cup Depend[M]) \setminus Losing$;
      else
        if $M \notin Processed$ **then**
          foreach $M' \in (Succs_{ctrl} \cup Succs_{env} \cup Succs_{delay})$ **do**
            if $M' \notin Losing \wedge M' \models \varphi_k$ **then**
              $Depend[M'] \leftarrow Depend[M'] \cup \{M\}$;
              $Waiting \leftarrow Waiting \cup \{M'\}$;
        $Processed \leftarrow Processed \cup \{M\}$;
  return *tt if* $cut(M_0) \notin Losing$, *else ff*

the controller cannot win from this marking, then M gets inserted into the set *Losing* while the dependencies of M are put to the set *Waiting* in order to backward propagate this information. If the initial marking is ever inserted to the set *Losing*, we can terminate and announce that a controller strategy does not exist. If this is not the case and there are no more markings in the set *Waiting*, then we terminate with success. In this case, it is also easy to construct the controller strategy by making choices so that the set *Losing* is avoided.

Theorem 3 (Correctness). *Algorithm 1 terminates and returns tt if and only if there is a controller strategy for the safety objective* $\varphi_k = \varphi \wedge k \geq \sum_{p \in P} p$.

1 Stream	$D = 133$	$D = 173$	$D = 213$	$D = 253$	$D = 293$	$D = 333$	$D = 373$
Tracks	70	90	110	130	150	170	190
TAPAAL	30.14s	69.78s	128.58s	216.44s	316.71s	491.65s	665.34s
UPPAAL	36.41s	76.63s	193.37s	351.17s	509.46s	1022.83s	1604.04s
2 Streams	$D = 19$	$D = 27$	$D = 35$	$D = 43$	$D = 51$	$D = 59$	$D = 67$
Tracks	6	8	10	12	14	16	18
TAPAAL	1.98s	7.34s	30.73s	101.92s	210.25s	398.00s	768.11s
UPPAAL	19.11s	93.46s	436.15s	1675.85s	3328.66s	◷	◷
3 Streams	$D = 17$	$D = 21$	$D = 25$	$D = 29$	$D = 35$	$D = 39$	$D = 43$
Tracks	3	4	5	6	7	8	9
TAPAAL	2.20s	16.52s	72.41s	244.28s	885.60s	(2132.71s)	◷
UPPAAL	885.56s	◷	◷	◷	◷	◷	◷

Table 1: Time in seconds to find a controller strategy for the disk operation scheduling for the smallest D where such a strategy exists.

6 Experiments

The discrete-time controller synthesis algorithm was implemented in the tool TAPAAL [15] and we evaluate the performance of the implementation by comparing it to UPPAAL TiGa [4] version 0.18, the state-of-the-art continuous-time model checker for timed games. The experiments were run on AMD Opteron 6376 processor limited to using 16 GB of RAM[1] and with one hour timeout (denoted by ◷).

6.1 Disk Operation Scheduling

In the disk operation scheduling model presented in Section 2 we scale the problem by changing the number of tracks and the number of simultaneous read streams. A similar model using the timed automata formalism was created for UPPAAL TiGa. We then ask whether a controller exists respecting a fixed deadline D for all requests. For each instance of the problem, we report the computation time for the smallest deadline D such that it is possible to synthesize a controller. Notice that the disk operating scheduling game net has an urgent controller, hence the discrete and continuous-time semantics coincide.

The results in Table 1 show that our algorithm scales considerably better than TiGa (that suffers from the large fragmentation of zone federations) as the number of tracks increases and it is significantly better when we add more read streams (and hence increase the concurrency and consequently the number of timed tokens/clocks).

[1] UPPAAL TiGa only exists in a 32 bit version, but for none of the tests the 4GB limit was exceeded for UPPAAL TiGa.

2 Processes/7-13 tokens

Max Age	10 Tasks		12 Tasks		14 Tasks		16 Tasks		18 Tasks	
5	(100)	**63s**	(**100**)	**141s**	(**100**)	**283s**	(**100**)	**570s**	(**100**)	**829s**
$D \leq 144$	(100)	100s	(98)	413s	(85)	1201s	(35)	◷	(18)	◷
10	(**100**)	318s	(**100**)	**882s**	(**96**)	**1555s**	(**65**)	**2911s**	(**14**)	◷
$D \leq 288$	(96)	**221s**	(69)	1443s	(43)	◷	(16)	◷	(1)	◷
15	(**99**)	1054s	(**78**)	2521s	(19)	◷	(**14**)	◷	(**2**)	◷
$D \leq 432$	(87)	**315s**	(60)	**1960s**	(19)	◷	(8)	◷	(0)	◷
20	(80)	2479s	(22)	◷	(14)	◷	(3)	◷	(**2**)	◷
$D \leq 576$	(**90**)	**554s**	(**66**)	**2914s**	(**34**)	◷	(**4**)	◷	(1)	◷

3 Processes/10-19 tokens

Max Age	2 Tasks		3 Tasks		4 Tasks		5 Tasks		6 Tasks	
5	(**100**)	**2s**	(**100**)	**39s**	(**99**)	**402s**	(**66**)	**1884s**	(**38**)	◷
$D \leq 57$	(99)	16s	(69)	1827s	(4)	◷	(0)	◷	(0)	◷
10	(**100**)	**15s**	(**97**)	**484s**	(**47**)	◷	(**20**)	◷	(**6**)	◷
$D \leq 114$	(98)	32s	(52)	3338s	(6)	◷	(0)	◷	(0)	◷
15	(**100**)	51s	(**69**)	**1373s**	(**28**)	◷	(**4**)	◷	(0)	◷
$D \leq 171$	(98)	**27s**	(50)	◷	(1)	◷	(0)	◷	(0)	◷

4 Processes/13-25 tokens

Max Age	2 Tasks		3 Tasks		4 Tasks		5 Tasks		6 Tasks	
5	(**92**)	**215s**	(**30**)	◷	(**7**)	◷	(**1**)	◷	(0)	◷
$D \leq 66$	(3)	◷	(0)	◷	(0)	◷	(0)	◷	(0)	◷
10	(**60**)	**2286s**	(**11**)	◷	(**2**)	◷	(0)	◷	(0)	◷
$D \leq 132$	(0)	◷	(0)	◷	(0)	◷	(0)	◷	(0)	◷

Table 2: Results for infinite scheduling of DPAs. The first row in each age-instance is TAPAAL, the second line is UPPAAL TiGa. The format is (X) Ys where X the number of solved instances (within 3600 seconds) out of 100 and Y is the median time needed to solve the problem. The largest possible constant for each row is given as an upper bound of the deadline D.

6.2 Infinite Job Shop Scheduling

In our second experiment, infinite job shop scheduling, we consider the duration probabilistic automata [29]. Kempf et al. [26] showed that "non-lazy" schedulers are sufficient to guarantee optimality in this class of automata. Here non-lazy means that the controller only chooses what to schedule at the moment when a running task has just finished (the time of this event is determined by the environment). We consider here a variant of this problem that should guarantee an infinite (cyclic) scheduling where all processes that share various resources and must meet their deadlines. The countdown of a process is started when its first task is initiated and the process deadline is met if the process is able to execute its last task within the deadline. After such a completed cycle, the process starts from its initial configuration and the deadline-clock is restarted. The task of the controller is now to find a schedule such that all processes always meet their deadline. The problem can be modelled using urgent controller, so the discrete and continuous-time semantics again coincide.

The problem is scaled by the number of parallel processes, number of tasks in each processes and the size of constants used in guards (excepted the deadline D that contains a considerably larger constant). For each set of scaling parameters, we generated 100 random instances of the problem and report on the number of cases where the tool answered the synthesis problem (within one hour deadline) and if more than 50 instances were solved, we also compute the median of the running time.

The comparison with UPPAAL TiGa in Table 2 shows a similar trend as in the previous experiment. Our algorithm scales nicely as we increase the number of tasks as well as the number of processes. This is due to the fact that the zone fragmentation in TiGa increases with the number of parallel components and more distinct guards. When scaling the size of constants, the performance of the discrete-time method gets worse and eventually UPPAAL TiGa can solve more instances.

7 Conclusion

We introduced timed-arc Petri net games and showed that for urgent controllers, the discrete and continuous-time semantics coincide. The presented discrete-time method for solving timed-arc Petri net games scales considerably better with the growing size of problems, compared to the existing symbolic methods. On the other hand, symbolic methods scale better with the size of the constants used in the model. In the future work, we may try to compensate for this drawback by using approximate techniques that "shrink" the constants to reasonable ranges while still providing conclusive answers in many cases, as demonstrated for pure reachability queries in [7]. Another future work includes the study of different synthesis objectives, as well as the generation of continuous-time strategies from discrete-time analysis techniques on the subclass of urgent controllers.

Acknowledgments. The research leading to these results has received funding from the EU FP7 FET projects CASSTING and SENSATION, the project DiCyPS funded by the Innovation Fund Denmark, the Sino Danish Research Center IDEA4CPS and the ERC Advanced Grant LASSO. The third author is partially affiliated with FI MU, Brno, Czech Republic.

References

1. Synt 2015. Electronic Proceedings in Theoretical Computer Science, 2015. http://formal.epfl.ch/synt/2015/.
2. R. Alur and D. L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, April 1994.
3. M. Andersen, H.G. Larsen, J. Srba, M.G. Sørensen, and J.H. Taankvist. Verification of liveness properties on closed timed-arc Petri nets. In *Mathematical and Engineering Methods in Computer Science: 8th International Doctoral Workshop*, volume 7721 of *LNCS*, pages 69–81. Springer, 2013.

4. G. Behrmann, A. Cougnard, A. David, E. Fleury, K. G. Larsen, and D. Lime. Uppaal-tiga: Time for playing games! In *Computer Aided Verification: 19th International Conference*, volume 4590 of *LNCS*, pages 121–125. Springer, 2007.
5. G. Behrmann, A. David, K.G. Larsen, J. Hakansson, P. Petterson, Wang Yi, and M. Hendriks. Uppaal 4.0. In *Third International Conference on Quantitative Evaluation of Systems*, pages 125–126, 2006.
6. B. Berthomieu and F. Vernadat. Time Petri nets analysis with TINA. In *Third International Conference on Quantitative Evaluation of Systems*, pages 123–124. IEEE Computer Society, 2006.
7. S.V. Birch, T.S. Jacobsen, J.J. Jensen, Ch. Moesgaard, N.N. Samuelsen, and J. Srba. Interval abstraction refinement for model checking of timed-arc Petri nets. In *Formal Modeling and Analysis of Timed Systems: 12th International Conference*, volume 8711 of *LNCS*, pages 237–251. Springer, 2014.
8. T. Bolognesi, F. Lucidi, and S. Trigila. From Timed Petri Nets to Timed LOTOS. In *Protocol Specification, Testing and Verification X, Proceedings of the IFIP WG6.1 Tenth International Symposium on Protocol Specification*, pages 395–408. North-Holland, 1990.
9. M. Bozga, C. Daws, O. Maler, A. Olivero, S. Tripakis, and S. Yovine. Kronos: A model-checking tool for real-time systems. In *Computer Aided Verification: 10th International Conference*, volume 1427 of *LNCS*, pages 546–550, 1998.
10. M. Bozga, O. Maler, and S. Tripakis. Efficient verification of timed automata using dense and discrete time semantics. In *Correct Hardware Design and Verification Methods: 10th IFIP WG10.5 Advanced Research Working Conference*, volume 1703 of *LNCS*, pages 125–141. Springer, 1999.
11. F. Cassez, A. David, E. Fleury, K. G. Larsen, and D. Lime. Efficient on-the-fly algorithms for the analysis of timed games. In *Concurrency Theory: 16th International Conference*, volume 3653 of *LNCS*, pages 66–80. Springer, 2005.
12. A. Church. Application of recursive arithmetic to the problem of circuit synthesis. *Journal of Symbolic Logic*, 28(4):289–290, 1963.
13. A. Church. Logic, arithmetic, and automata. In *Proc. Internat. Congr. Mathematicians (Stockholm, 1962)*, pages 23–35. Inst. Mittag-Leffler, 1963.
14. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. Introduction to algorithms. third edition., 2009.
15. A. David, L. Jacobsen, M. Jacobsen, K.Y. Jørgensen, M.H. Møller, and J. Srba. TAPAAL 2.0: Integrated development environment for timed-arc Petri nets. In *Tools and Algorithms for the Construction and Analysis of Systems: 18th International Conference*, volume 7214 of *LNCS*, pages 492–497. Springer, 2012.
16. D.L. Dill. Timing assumptions and verification of finite-state concurrent systems. In *Automatic Verification Methods for Finite State Systems: International Workshop*, volume 407 of *LNCS*, pages 197–212. Springer, 1990.
17. B. Finkbeiner. Bounded synthesis for Petri games. In *Correct System Design: Symposium in Honor of Ernst-Rüdiger Olderog on the Occasion of His 60th Birthday*, volume 9360 of *LNCS*, pages 223–237. Springer, 2015.
18. B. Finkbeiner and E. Olderog. Petri games: Synthesis of distributed systems with causal memory. In *Proceedings Fifth International Symposium on Games, Automata, Logics and Formal Verification*, volume 161 of *EPTCS*, pages 217–230, 2014.
19. B. Finkbeiner and H. Peter. Template-based controller synthesis for timed systems. In *Tools and Algorithms for the Construction and Analysis of Systems: 18th International Conference*, volume 7214 of *LNCS*, pages 392–406. Springer, 2012.

20. D. Gale and F. M. Stewart. Infinite games with perfect information. In *Contributions to the Theory of Games, Vol. 2*, Annals of Mathematics Studies, no. 28, pages 245–266. Princeton University Press, Princeton, N. J., 1953.
21. G. Gardey, D. Lime, M. Magnin, and O. Roux. Romeo: A tool for analyzing time Petri nets. In *Computer Aided Verification: 17th International Conference*, volume 3576 of *LNCS*, pages 261–272. Springer, 2005.
22. Y. Gurevich. Games people play. In *The Collected Works of J. Richard Büchi*, pages 517–524. Springer, 1990.
23. H. Hanisch. Analysis of Place/Transition Nets with Timed Arcs and its Application to Batch Process Control. In *Application and Theory of Petri Nets 1993: 14th International Conference*, volume 691 of *LNCS*, pages 282–299. Springer, 1993.
24. P. G. Jensen, K. G. Larsen, J. Srba, M. G. Sørensen, and J. H. Taankvist. Memory efficient data structures for explicit verification of timed systems. In *NASA Formal Methods: 6th International Symposium*, volume 8430 of *LNCS*, pages 307–312. Springer, 2014.
25. K.Y. Jørgensen, K. G. Larsen, and J. Srba. Time-darts: A data structure for verification of closed timed automata. In *Proceedings Seventh Conference on Systems Software Verification*, volume 102 of *EPTCS*, pages 141–155. Open Publishing Association, 2012.
26. J-F. Kempf, M. Bozga, and O. Maler. As soon as probable: Optimal scheduling under stochastic uncertainty. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, volume 7795 of *LNCS*, pages 385–400. Springer, 2013.
27. K. G. Larsen and Y. Wang. Time-abstracted bisimulation: Implicit specifications and decidability. *Information and Computation*, 134(2):75 – 101, 1997.
28. X. Liu and S. A. Smolka. Simple linear-time algorithms for minimal fixed points (extended abstract). In *Automata, Languages and Programming: 25th International Colloquium*, volume 1443 of *LNCS*, pages 53–66. Springer, 1998.
29. O. Maler, K. G. Larsen, and B. H. Krogh. On zone-based analysis of duration probabilistic automata. In *Proceedings 12th International Workshop on Verification of Infinite-State Systems*, volume 39 of *EPTCS*, pages 33–46. 2010.
30. J.A. Mateo, J. Srba, and M.G. Sørensen. Soundness of timed-arc workflow nets in discrete and continuous-time semantics. *Fundamenta Informaticae*, 140(1):89–121, 2015.
31. H. Peter. Component-based abstraction refinement for timed controller synthesis. pages 364–374. IEEE, IEEE Computer Society, 2009.
32. H. Peter, R. Ehlers, and . Mattmüller. Synthia: Verification and synthesis for timed automata. In *Computer Aided Verification: 23rd International Conference*, volume 7214 of *LNCS*, pages 649–655. Springer, 2011.
33. A Pnueli, E Asarin, O Maler, and J Sifakis. Controller synthesis for timed automata. In *System Structure and Control*. Citeseer, Elsevier, 1998.
34. J.F. Raskin, M. Samuelides, and L.V. Begin. Petri games are monotone but difficult to decide. Technical report, Université Libre De Bruxelles, 2003.
35. V.V. Ruiz, F. Cuartero Gomez, and D. de Frutos Escrig. On non-decidability of reachability for timed-arc Petri nets. In *The 8th International Workshop on Petri Nets and Performance Models. Proceedings.*, pages 188–196, 1999.
36. Qiong Zhou, Michael Wang, and S.P. Dutta. Generation of optimal control policy for flexible manufacturing cells: A Petri net approach. *The International Journal of Advanced Manufacturing Technology*, 10(1):59–65, 1995.

A Proof of Theorem 1

We initially define a witness for the nonexistence of a strategy for a given TAPG G with urgent controller. The intuition of the witness is to provide a strategy for the environment s.t. it describes all possible choices of the controller. Thus, for the environment choices there are only singleton continuations (if any), and for controller choices the multitude of possible continuations.

Definition 7. *A witness is a function $\gamma : \Pi_G \to 2^{\mathcal{M}(G)}$ for a marked TAPG G with urgent controller, which for a sequence of markings π defines the next possible markings in the sequence for the environment. A function γ must satisfy the following conditions for every π (here Last(π) is the last marking in the sequence π):*

- *If Last(π) $\not\models \varphi$ then $\gamma(\pi) = \emptyset$.*
- *Else if there is no $d \in \mathbb{R}^{\geq 0}$ and no $t \in T$ s.t. Last(π) $\xrightarrow{d}$ $M' \xrightarrow{t}$ then $\gamma(\pi) = \emptyset$.*
- *Else if for all $t \in T_{ctrl}$ it holds that Last(π) $\overset{t}{\not\to}$ then $\gamma(\pi) = \{M''\}$ where for some $d \in \mathbb{R}^{\geq 0}$ and some $t \in T_{env}$ it holds that Last(π) $\xrightarrow{d}$ $M' \xrightarrow{t} M''$.*
- *Else*
  - *either $\gamma(\pi) = \{M''\}$ s.t. Last(π) $\xrightarrow{t} M''$ for some $t \in T_{env}$ or*
  - *$\gamma(\pi) = \{M'' \mid \text{Last}(\pi) \xrightarrow{t} M'', t \in T_{ctrl}\}$.*

We write $\Gamma_0 = \{M_0\}$ for all sequences of length 1 from the initial marking M_0, and $\Gamma_n = \bigcup_{\pi \in \Gamma_{n-1}} \pi \circ \gamma(\pi)$ for all continuations of lengths less than $n+1$.

A witness γ disproves the existence of any strategy ensuring *Safe*(φ), if $\exists n \in \mathbb{N}$ s.t. $\Gamma_n = \Gamma_{n+1}$, implying that Γ_n describes a finite tree, and $\forall \pi \in \Gamma_n$ there exists $M \in \pi$ s.t. $M \not\models \varphi$ showing that every branch of the tree eventually breaks φ. We call such a witness a counter witness.

Given a counter witness disproving the existence of any strategy, we can construct a linear constraint-system describing a family of witnesses. This technique is an adaption of the one used by Mateo et al. [30].

Let G be a marked TAPG with urgent controller, and γ be a counter witness of φ, then we can view γ as a tree as illustrated in Figure 3. Let us dwell shortly on the indexing used throughout the text. For indexing we here use sequences, and write $a = a_1 a_2 \dots a_n$ for a sequence. We define the minus operator $a - k = a_1 \dots a_{n-k}$ and implicit concatenation $ab = a_1 \dots a_n b_1 b_2 \dots$. We say that $a \lhd b$ if a is a prefix of b and $a \neq b$. We also define a total ordering of indexes, s.t. if $\forall n < k$ it holds that $a_n = b_n$ and $a_k < b_k$ then $a < b$. In the following, it suffices to use single-digits for each element of an index. For all labels in γ we write $\mathcal{L}_\gamma$ and omit γ when it is obvious from the context. The intuition is that if $a \lhd b$ then M_a is in the path leading to M_b.

The idea is to create a constraint system from the witness where the concrete delays are replaced by variables, and then solve the resulting constraint system. Let us now construct a table Θ, assisting us in creating this constraint system.

Formally, a table Θ for the witness γ is a matrix with m rows and n columns such that each element $\Theta_{y,a}$ of the table, where $1 \leq y \leq m$ and $a \in \mathcal{L}$, contains

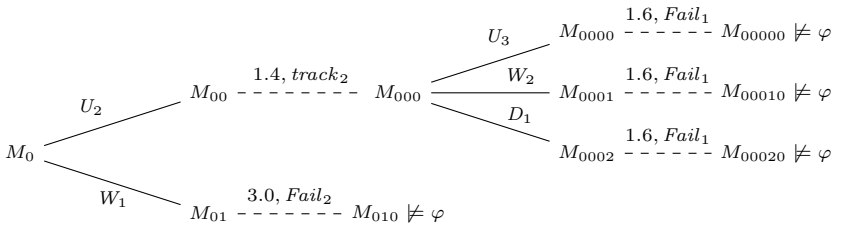

Fig. 3: The tree representation of a counter witness for $\varphi = Fail \leq 0$ for the example in Figure 1 when $D = 3$. The labels of the branches show which place will receive a new token. Dashed lines indicate environment choices.

either the value $\perp$ (unused token) or the pair (p, f) where $p \in P$ represents the location of the token and $f \in \mathbb{N}_0 \cup \{\bullet\}$ is a flag signalling whether the age of the given token was set to some given value[2] from $\mathbb{N}_0$ or whether it has not changed (the flag value $\bullet$). We denote by $\Theta^{place}_{y,a}$ and $\Theta^{flag}_{y,a}$ the elements p and f of the pair of $\Theta_{y,a}$, respectively. If one element of the pair (p, f) is not relevant for our considerations, we simply write $(p, -)$ or $(-, f)$. We let y to range over the rows in the table (tokens) and a over the columns in the table (transition firing steps).

Definition 8 (Valid table for witness γ). *A table Θ for witness γ is* valid *if the following conditions are met.*

a) *Given the initial marking $M_0 = \{(p_1, x_1), (p_2, x_2), \ldots, (p_k, x_k)\}$, the zero column of Θ is defined as $\Theta_{y,0} \stackrel{def}{=} (p_y, x_y)$ if $1 \leq y \leq k$, and $\Theta_{y,0} \stackrel{def}{=} \perp$ if $k < y \leq m$.*

b) *For each column a in the table where $M_a \models \varphi$ there exists a set of transitions $\{t_{a0}, \ldots t_{ab} \ldots, t_{ac}\}$ s.t. $\exists d \in \mathbb{R}^{\geq 0}$ s.t. $M_a \stackrel{d}{\rightarrow} M'_a \stackrel{t_{ab}}{\rightarrow} M_{ab}$, which gives us the two sets $Consume_{ab}$, $Produce_{ab}$ and the bijection $\mathcal{U}_{ab} : \{0, \ldots, m\} \rightarrow \{0, \ldots, m\}$ for which it holds:*

- *The set $Consume_{ab}$ represents the y-indexes of the tokens in column a consumed by firing the transition t_{ab} such that for all $p \in {}^\bullet t_{ab}$ we have $w(p, t_{ab}) = |\{y \in Consume_{ab} \mid \Theta^{place}_{y,a} = p\}|$, and for all $p \in P \setminus {}^\bullet t_{ab}$ we have $\{y \in Consume_{ab} \mid \Theta^{place}_{y,a} = p\} = \emptyset$,*
- *The set $Produce_{ab}$ represents the y-indexes of the tokens in column ab produced by firing the transition t_{ab} such that for all $p \in t^\bullet_{ab}$ we have $w(t_{ab}, p) = |\{y \in Produce_{ab} \mid \Theta^{place}_{y,ab} = p\}|$ and for all $p \in P \setminus t^\bullet_{ab}$ we have $\{y \in Produce_{ab} \mid \Theta^{place}_{y,ab} = p\} = \emptyset$, and*
- *there is a bijection $\mathcal{U}_{ab} : \{1, \ldots, m\} \rightarrow \{1, \ldots, m\}$ that relates the indexes of column a with those presented in column ab such that*
  1. *if $|Consume_{ab}| \leq |Produce_{ab}|$ and $y \in Consume_{ab}$ then $\mathcal{U}_{ab}(y) \in Produce_{ab}$,*

[2] The age value can only be reset to 0.

2. *if* $|Consume_{ab}| \geq |Produce_{ab}|$ *and* $\mathcal{U}_{ab}(y) \in Produce_{ab}$ *then* $y \in Consume_{ab}$,
3. *if* $y \in Consume_{ab}$ *and* $Type((\Theta^{place}_{y,a}, t_{ab})) = Transport_j = Type((t_i, p'))$ *then* $\mathcal{U}_{ab}(y) = y$ *and* $\Theta^{place}_{y,ab} = p'$,
4. *if* $y \in \{1, \ldots, m\} \setminus Consume_{ab}$ *and* $\Theta_{y,a} \neq \perp$ *then* $\mathcal{U}_{ab}(y) = y$ *and* $\Theta^{place}_{y,ab} = \Theta^{place}_{y,a}$,
5. *if* $y \in \{1, \ldots, m\} \setminus Consume_{ab}$ *and* $\Theta^{place}_{y,a} = \perp$ *then either* $\mathcal{U}_{ab}(y) \in Produce_{ab}$, *or* $\mathcal{U}_{ab}(y) = y$ *and* $\Theta^{place}_{y,ab} = \perp$, *and*
6. *if* $\Theta_{\mathcal{U}_{ab}(y),ab} = \perp$ *then* $y \in Consume_{ab}$ *or* $\Theta^{place}_{y,ab} = \perp$.

c) *For each column* ab *in the table:*

- *if* $Type((p, t_{ab})) = Inhib$ *for some* $p \in P$ *then* $|\{y \in \{1, \ldots, m\} \mid \Theta^{place}_{y,a} = p\}| < w(p, t_{ab})$
- *for all* $y \in Produce_{ab}$, *if* $Type((t_{ab}, \Theta^{place}_{y,ab})) = Normal$ *then* $T^{flag}_{y,ab} = 0$ *else* $T^{flag}_{y,ab} = \bullet$, *and*
- *if* $y \notin Produce_{ab}$ *and* $\Theta_{y,ab} \neq \perp$ *then* $T^{flag}_{y,ab} = \bullet$.

We can now transform the tree from Figure 3, into such a table, as presented in Table 3.

	M_0	M_{00}	M_{000}	M_{0000}	M_{00000}	M_{0001}
1	$(R_1, 0)$	$(R_1, \bullet)$	$(R_1, \bullet)$	$(R_1, \bullet)$	$(Fail_1, 0)$	$(R_1, \bullet)$
2	$(R_2, 0)$	$(R_2, \bullet)$	$(R_2, \bullet)$	$(R_2, \bullet)$	$(R_2, \bullet)$	$(W_2, 0)$
3	$(R_3, 0)$	$(R_3, \bullet)$	$(R_3, \bullet)$	$(R_3, \bullet)$	$(R_3, \bullet)$	$(R_3, \bullet)$
4	$(track_1, 0)$	$(U_2, 0)$	$(track_2, 0)$	$(U_3, 0)$	$(U_3, \bullet)$	$\perp$

	M_{00010}	M_{0002}	M_{00020}	M_{01}	M_{010}
1	$(Fail_1, 0)$	$(R_1, \bullet)$	$(Fail_1, 0)$	$(W_1, 0)$	$(W_1, \bullet)$
2	$(W_2, \bullet)$	$(R_2, \bullet)$	$(R_2, \bullet)$	$(R_2, \bullet)$	$(Fail_2, 0)$
3	$(R_3, \bullet)$	$(R_3, \bullet)$	$(R_3, \bullet)$	$(R_3, \bullet)$	$(R_3, \bullet)$
4	$\perp$	$(D_1, 0)$	$(D_1, 0)$	$\perp$	$\perp$

Table 3: A valid table for the witness in Figure 3.

One can verify that there is a valid table for any computation of the given TAPG G. On the other hand, any valid table defines a legal computation on untimed markings represented by each column.

Definition 9 (Untimed marking given by column a). *Let Θ be a valid table and let $a \in \mathcal{L}$. We define the untimed marking $M^u_a \stackrel{def}{=} \{\Theta^{place}_{y,a} \in P \mid 1 \leq y \leq m\}$ as a multiset of all places where a token is present in the column a of the table T.*

By the way in which a valid table is constructed, we can verify the validity of the following lemma.

Lemma 4 (Untimed consistency of a valid table Θ). *Let Θ be a valid table. Then $M^u_a \stackrel{t_{ab}}{\rightarrow} M^u_{ab}$ for all $a, ab \in \mathcal{L}$, in the classical (untimed) Petri net semantics.*

We shall now proceed with defining a set of difference constraint inequalities that encode in a sound and complete way the timings aspects of the witness depicted in Figure 3.

Let the *execution time* of a transition t_{ab} in Figure 3 be denoted by the variable e_{ab} representing the total time elapsed from the initialization until the transition t_{ab} is fired. In order to construct the system of inequalities over the variables $e_a, \ldots, e_{ab}$, we need to define an expression describing the age of a token y just at the moment when the transition t_{ab} is fired.

Definition 10 (Token-age expression). *Let Θ be a valid table. We define $age(y, b)$, where $1 \leq y \leq m$ and $ab \in \mathcal{L}$, as the expression*

$$\text{“}e_{ab} - e_a + d\text{”}$$

a is a prefix of a', $|a| + 1 = |a'|$ and a' is a prefix of ab where a' is the longest prefix, such that $\Theta^{flag}_{y,a} = d \in \mathbb{N}_0$. By agreement, e_ϵ is replaced with 0.

The intuition is that $age(y, ab)$ expresses, in terms of the execution time variables, the current age of the token y after the time delay d_{ab} and just before firing the transition t_{ab}. The correctness of the definition follows from the requirements on a valid table and the fact that the first column ($ab = \epsilon$) of any valid table contains only pairs (p, f) where $f \neq \bullet$.

We are now ready to define, for a given valid table Θ, a system of inequalities over the variables e_a for all $a \in \mathcal{L}$ that expresses all the timing constraints on the firing of transitions, age invariants and urgency.

Definition 11 (Constraint system). *Let Θ be a valid table for a witness γ from depicted in Figure 3. The constraint system $\mathcal{C}$ for Θ is the set of inequations over the variables $\ldots, e_a, \ldots, e_c, \ldots$, where for each label $a, c \in \mathcal{L}$ where $a \lhd c$, then $e_a \leq e_c$. We create further inequalities, such that for all $p \in P$, $y \in \{1, \ldots, m\}$ and $a \in \mathcal{L}$:*

a) *if $\Theta^{place}_{y,a} = p$ and $I(p) = [0, u]$ where $u \in \mathbb{N}_0$, we add the inequality $age(y, a) \leq u$ to $\mathcal{C}$,*
b) *if $\Theta^{place}_{y,a} = p$ and $y \in Consume_a$ and $(p, t_a) \in IA$ and $g((p, t_a)) = [l, u]$, we add $l \leq age(y, a)$ and if $u \neq \infty$ also $age(y, a) \leq u$ to $\mathcal{C}$, and*
c) *if M_a enables some $t \in T_{urg}$ then we add $e_a - e_{a-1} = 0$ to $\mathcal{C}$ where, as mentioned above, e_ϵ is replaced with 0.*

In our witness depicted in Figure 3, the constraint system for the valid table shown in Table 3 is the following.

– We add the inequalities

$$e_0 \leq e_{00}, e_{00} \leq e_{000}, e_{000} \leq e_{0000}, e_{0000} \leq e_{00000}.$$

$$e_{000} \leq e_{0001}, e_{0001} \leq e_{00010}.$$

$$e_{000} \leq e_{0002}, e_{0002} \leq e_{00020}.$$

$$e_0 \leq e_{01}, e_{01} \leq e_{010}.$$

– For the nontrivial age invariants, we add

$$age(2, {}_{0001}) \leq 4, age(2, {}_{00010}) \leq 4, age(1, {}_{01}) \leq 4, age(1, {}_{010}) \leq 4$$
$$age(4, {}_{00}) \leq 2, age(4, {}_{0000}) \leq 2, age(4, {}_{00000}) \leq 2, age(4, {}_{0002}) \leq 2$$
$$age(4, {}_{00020}) \leq 2$$

– Regarding urgency, we add the constraints

$$e_0 = 0, e_{000} = e_{00}$$

– Finally for the the guards on input arcs, we add the constraints

$$3 \leq e_{020} - e_{02} \leq 3, 3 \leq e_{00020} - e_{0002} \leq 3, 3 \leq e_{00010} - e_{0001} \leq 3$$
$$3 \leq e_{00000} - e_{0000} \leq 3, 1 \leq e_{00} - e_0 \leq 2, 1 \leq e_{0000} - e_{000} \leq 2$$
$$1 \leq e_{0002} - e_{000} \leq 2, 1 \leq e_{0001} - e_{000} \leq 4, 1 \leq e_{01} - e_0 \leq 4$$

Observe that the original delays in the witness from Figure (3) form a solution of the constructed constraint system: $e_0 = 0$, $e_{00} = 1.4$, $e_{000} = 1.4$, $e_{0000} = 3.0$, $e_{00000} = 3.0$, $e_{0001} = 3.0$, $e_{00010} = 3.0$, $e_{0002} = 3.0$, $e_{00020} = 3.0$, $e_{01} = 3.0$, $e_{010} = 3.0$ In fact, there is also an integer solution to the constraint system (this is not only a coincidence), e.g. $e_0 = 0$, $e_{00} = 2$, $e_{000} = 2$, $e_{0000} = 3$, $e_{00000} = 3$, $e_{0001} = 3$, $e_{00010} = 3$, $e_{0002} = 3$, $e_{00020} = 3$, $e_{01} = 3$, $e_{010} = 3$ and such a tree is also a counter witness in our TAPG with urgent controller.

Lemma 5. *Let γ be counter witness for a TAPG G. Then there is a valid table Θ for γ and the corresponding constraint system $\mathcal{C}$ such that $e_c = \sum_{\forall a \in \mathcal{L},\, a \lhd c} d_c$ is a solution of $\mathcal{C}$. Moreover, $e_0, \ldots, e_c$ is a (real) solution of $\mathcal{C}$ if and only if $M_0 \overset{e_a, t_a}{\to} M_a \overset{e_{ab}, t_{ab}}{\to} \ldots \overset{e_c, t_c}{\to} M_c$*

Proof. By analyzing the requirements for a valid table, we can see that if all the sequences $\pi \in \Gamma$ defined by a witness γ can be performed in the TAPG G then we are able to design a table that satisfies all the requirements for the untimed part of the sequence execution and uses the right tokens in the pairing bijection such that the corresponding constraint system $\mathcal{C}$ gives the sufficient and necessary conditions for the execution time variables e_a to produce a valid computation of the TAPG G with the given sequence of transitions firing. Hence the original execution times in the given sequence form one possible solution of the system but any such a solution actually provides a possible timed execution of the sequence (note that if a transition t_a is performed at time e_a and transition t_{ab} is performed at time e_{ab} then the delay between the execution of these two transitions is $e_{ab} - e_a$).

We can now summarise and conclude the proof of Theorem 1. We assumed a witness for the game with real delays as in Figure 3. Based on this we know that there exists a valid table Θ for such a witness such that the constraint system $\mathcal{C}$

for the witness, representing all possible delays that can execute the transition sequences in Figure 3, has a solution corresponding to the delays in Figure 3. This is due to Lemma 5.

The constraint system $\mathcal{C}$ is an instance of linear programming problem where we used difference constraints only. We can therefore reduce the problem to a shortest-path problem with integer weights only, this implies that an integer solution exists [14]. Hence using Lemma 5 we know that there is also a witness disproving φ in G following the same transitions as in Figure 3 but with integral delays only.

B Proof of Theorem 3

The **while**-loop of the algorithm presented in 1 has the following invariant when running on the TAPG G;

Lemma 6 (Loop invariant). *If $M \in$ Losing then there exists no winning strategy from M satisfying Safe(φ_k).*

We further state the finiteness of the internal data-structures.

Lemma 7. *Waiting, Processed and Losing are all finite.*

Initially we can observe that markings added to *Processed* and *Losing* all are popped from *Waiting*, therefor it is sufficient to show that only a finite set of markings can be added to *Waiting*. Given that we work only with discrete time delays and under the extrapolation of *cut*, we know that for a given distribution of tokens, there exists only finitely many unique markings. Given some bound k, limiting the number of tokens in the game, then it follows that the permutations of token-distribution is finite, and hence the set of markings going into *Waiting* must also be so. From line 2 in the algorithm, we can see that φ_k is defined such that $M \models \varphi_k$ only if M has less than k tokens. We can see on line 23, that only if a marking respects this property, it is added to the *Waiting* and *Depend* sets, and thus, *Waiting* is finite throughout the algorithm.

Given Lemma 7 we can now prove termination of Algorithm 1.

Proof (Termination). In each iteration of the algorithm, *Waiting* is initially decreased in size by 1. Only if we cannot directly determine if the marking M popped from *Waiting* is losing or not, we add more markings to *Waiting*. One of four things will occur:

1. If $M \notin$ *Processed* and M is not determined to be losing, then we add all successors M' of M to *Waiting* and *Depend*$[M]$, where $M' \models \varphi_k$ and $M' \notin$ *Losing*.
2. If $M \in$ *Processed*, and M is determined to be losing, then *Depend*$[M]$ is added to waiting, and all *Losing* are removed from waiting.
3. If $M \notin$ *Processed* and M is determined to be losing, then *Depend*$[M]$ is empty, and no new markings are added to *Waiting*.

4. If $M \in \mathit{Processed}$ and M is not determined to be losing, then nothing is added to $\mathit{Waiting}$.

From this it follows that if $M \in \mathit{Waiting}$, then it is so because either (1) $M \notin \mathit{Processed}$, or (2) some successor M' of M was determined to be losing. From line 26 it is trivial to verify that (1) only occurs once per marking. Given that we know that $\mathit{Losing} \cap \mathit{Waiting} = \emptyset$ (lines 15, 19, 23), then we know that (2) also only occurs finitely often per marking. Combined with the finiteness of $\mathit{Waiting}$ and $\mathit{Succs}[M]_{\{env,ctrl,delay\}}$ from Lemma 7 we have that the algorithm terminates. □

Soundness

We prove Soundness by proving that if $M \in \mathit{Losing}$ then there exists no winning strategy from M satisfying $\mathit{Safe}(\varphi_k)$. We then show that this holds inductively, and so if $\mathit{cut}(M_0) \in \mathit{Losing}$ then no strategy can exist.

Proof (Soundness). Our induction hypothesis is that for all M in Losing, no strategy exists from M satisfying $\mathit{Safe}(\varphi_k)$. Initially our assumption holds as $\mathit{Losing} = \emptyset$. Now let us investigate the behaviour when we add markings to Losing. Let us consider the three cases where we add markings to Losing.

- (Line 14 (a)) If there exists a $t \in T_{env}$ for the environment to a marking M' s.t. $M' \not\models \varphi_k$, M is also losing as the controller cannot defend against this "attack".
- (Line 14 (b)) If there exists a $t \in T_{env}$ to a marking $M' \in \mathit{Losing}$, then the environment can force the game into a setting where no strategy exists, the controller cannot defend.
- (Line 18) If the controller has some choice, but all the choices of the controller leads to markings $M' \in \mathit{Losing}$, then the controller is forced to "choose" badly.

As all three cases respect our induction hypothesis, we can conclude that the algorithm is sound. □

Completeness

We now continue to prove completeness. Intuitively we assume that the algorithm erroneously returns tt, which implies that it does not fully discover some play leading to a marking violating φ_k.

Proof (Completeness). Let us by contradiction assume that for some TAPG G the algorithm returns tt but there exists no strategy ensuring $\mathit{Safe}(\varphi_k)$, then surely one of three things must have occurred.

1. There exists a marking M^1, where $M^1 \not\models \varphi_k$, which was never discovered.
2. Upon termination there exists some $M^2 \in \mathit{Processed}$ where $\exists t \in T_{env}$ s.t. $M^2 \xrightarrow{t} M'$ and $M' \in \mathit{Losing}$ or $M' \not\models \varphi_k$.

3. Upon termination there exists some $M^3 \in \mathit{Processed}$ where $\forall \alpha \in T_{ctrl} \cup \{1\}$ it holds that $M^3 \xrightarrow{\alpha} M'$ and $M' \in \mathit{Losing}$ or $M' \not\models \varphi_k$.

As we are assuming the algorithm is falsely returning tt, this implies that *Waiting* must be empty upon termination. For case 2 and 3, it is easy to verify that if we in one step can reach/ are forced to reach a losing marking, or a marking invalidating $\mathit{Safe}(\varphi_k)$, then the algorithm detects this (lines 13 and 17) and backwards propagate the information by adding dependencies (which are all the immediate ancestors of the given marking which were ever on *Waiting*, which could not be determined to be losing – line 24) to *Waiting* (lines 15, 19). This violates our contradiction.

We now focus on case 1 and assume *wlog* that there exists only one M^1 where $M^1 \not\models \varphi_k$. Given lines 10, 11 and 12, we know that all successors are produced for any marking which has ever been added to *Waiting*, we also know that a finite sequence $\pi = M_0 M_1 \dots M_k \dots M_n M^1$ exist. Then either M_n is eventually added to *Losing*, which violates our contradiction, or for some $0 \leq k < n$ it holds that $M_k \in \mathit{Losing}$ in which case we are in cases 2 or 3, implying that we already found M_n. $\square$