Reaching as Cheap as Possible in 1-Clock Robust Weighted Timed Games

Nathalie Bertrand
Univ Rennes, Inria, CNRS, IRISA, France

Maëlle Gautrin
ENS Paris-Saclay, France

Julie Parreaux
Univ Rennes, Inria, CNRS, IRISA, France

Abstract

The value problem for 2-player games on graph generally consists in determining the minimal value Min can ensure against any possible strategy for Max. We consider here the value problem for reachability objectives in weighted timed games (WTGs) under a robust semantics. WTGs are a modelling formalism combining real-time constraints and integer weights on transitions and locations in an adversarial setting. Robustness allows for representing timing imprecisions in the measurement of delays and clock values. Robust weighted timed games have been introduced more than a decade ago: they are undecidable in general, and were quite recently shown decidable for the subclasses of acyclic or divergent robust WTGs. This paper pursues the goal of identifying decidable subclasses and establishes the decidability of the robust value problem for 1-clock WTGs.

2012 ACM Subject Classification Theory of computation → Timed and hybrid models; Theory of computation → Automata over infinite objects; Theory of computation → Logic and verification

Keywords and phrases timed automata, weighted timed games, robustness, games on graphs

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2026.19

Related Version *Full Version*: `https://arxiv.org/abs/2606.28773` [3]

Funding This work has been partially funded by the project ANR BisoUS (ANR-22-CE48-0012).

1 Introduction

Timed Automata and Weighted Timed Games

imed automata form an elegant and well-studied model for real-time systems [2]. To check for the existence of an error execution, which can be phrased as a reachability question, efficient algorithms and tools have been developed since then.

To represent real-time controller synthesis problems, extensions of timed automata with weights and two antagonistic players have been introduced under the terminology of *weighted timed games* (WTGs, for short) [1, 4]. In WTGs, locations and transitions are equipped with integer weights. During an execution,transition weights are summed and added to the location weights scaled by the elapsed time. The objective of Player Min is to reach a target set of locations while minimising the cumulative weight. The value problem asks whether Min can ensure the cumulative weight is below a given threshold.

Interestingly, the presence of both negative and positive weights makes the value problem for WTGs challenging: it is undecidable in general [11]. This negative result for WTGs with at least three clocks has been recently improved to show the undecidability of the value problem for WTGs with two clocks [18]. When restricting to 1-clock WTGs, i.e. WTGs for which the underlying timed automaton has a single clock, it has been shown that the value can be computed by various algorithms [6, 19], yet the value problem is PSPACE-hard [15].

The decidability frontier of the related problem of approximating the value in WTGs has also been studied. It is undecidable in general [17] and decidable for restricted classes for non-negative weights only [5] and also including negative weights [13].

Robustness

Timed automata are an elegant model, but they expose some undesirable artefacts from mathematical modelling. Specifically, the precision in time measurement is supposed to be arbitrary, which is unrealistic for practical real-time systems. Towards implementability, robustness has been considered in the literature. A parametric semantics that allows some perturbation of delays (up to a parameter bound) was introduced in [25]. The early works consider a robust semantics for timed automata, in which perturbations are controlled by an opponent, while the controller has a reachability objective [27, 28, 8, 7]. It thus amounts to deciding whether there exists a robust execution reaching a target set of locations. To solve this problem, the data structure of Difference Bound Matrix (DBM), which is standard in algorithms for timed automata, has been refined to shrunk DBMs [26]. While several variants of robustness have been considered in the literature, we will focus on the conservative semantics [27, 24, 14]. There, the delays proposed by the controller before firing a transition must remain feasible under any perturbation (within the allowed perturbation bound) proposed by its opponent.

In the context of weighted timed games, the perturbations are naturally in the hands of the opponent Max, who is antagonistic to Min. The robust value problem asks whether Min can guarantee reaching the target set of locations with the cumulative weight bounded by a threshold. Robust strategies must therefore anticipate any delay perturbation Max could make. The robust value is then defined for infinitesimal perturbations, as in the timed automata's case [25, 27]. The decidability landscape for the robust value problem is currently as follows. Perhaps unsurprisingly, like for non-robust WTGs, it was shown a decade ago that the robust value problem is undecidable in general [8, 16]. More recently, classes of WTGs for which the robust value problem is decidable have been identified, notably divergent WTGs (introduced in [5, 12]) and acyclic WTGs [22]. For acyclic WTGs, the decidability proof can be seen as an involved parametric extension of the fixpoint computation of the (exact) value [1]. For divergent WTGs, a finite unfolding of the game is performed to return to acyclic WTGs. In the special case of 1-clock WTGs, to our knowledge, the only decidability result concerns restricted 1-clock WTGs in which all weights are non-negative, and with an assumption about the cycles [16]. The proof technique generalises the construction for 1-clock WTGs under the classical (i.e. non-robust) semantics [7].

Contribution

In this paper, we establish the decidability of the robust value problem for a wide class of 1-clock WTGs with arbitrary integer weights. The only assumptions we make are common when studying robust WTGs, see [22] for instance: first the absence of states with exact value $-\infty$, and second the clock valuation is bounded by a maximal constant M.

To prove the decidability of the robust value problem for 1-clock WTGs, we transform the original WTG $\mathcal{G}$ into another WTG $\mathcal{C}$ formed of several copies of $\mathcal{G}$. Intuitively, a copy indexed by some integer a represents that the clock valuation is bounded below by a. This accounts for the fact that Max can use delay perturbations to enforce lower-bounds on the clock valuation (until the potential next reset). We prove that exact value in the copy game $\mathcal{C}$ coincides with the robust value in the original game $\mathcal{G}$. Since the value problem for 1-clock

WTGs is decidable [23], we obtain the desired decidability result. A fundamental property of the copy game is that it exhibits no deadlocks caused by robustness. Therefore the exact value and robust value coincide in $\mathcal{C}$. A key in our development is thus to identify deadlocks that are not present in the classical exact semantics and appear in the robust semantics.

Due to space constraints, all omitted proofs can be found in the related long version [3].

2 Robust Weighted Timed Games

We consider weighted timed games with a single *clock*, denoted by x. The *valuation* of x is a non-negative real number $\nu \in \mathbb{R}_{\geq 0}$. For a valuation ν, a delay $t \in \mathbb{R}_{\geq 0}$ and a set $Y \subseteq \{x\}$, we uniformly write $\nu[Y := 0]$ for the valuation obtained by a reset of Y, that is $\nu[Y := 0] = \nu$ if $Y = \emptyset$ and $\nu[Y := 0] = 0$ if $Y = \{x\}$. A *guard* is denoted by an interval $I \subseteq \mathbb{R}^+$ of non-negative real values with bounds in $\mathbb{N} \cup \{+\infty\}$. It constrains the valuations of the clock: a valuation ν satisfies a guard I, denoted $\nu \models I$, if $\nu \in I$. Given a guard I, we denote by $\overline{I}$ (resp. $\mathsf{left}(I)$, and $\mathsf{right}(I)$) the closure (resp. lower, and upper bounds) of I. We let Guards denote the set of guards over the clock x.

We now introduce the central notion of this paper, namely weighted timed games.

▶ **Definition 1.** *A* weighted timed game *(WTG) is a tuple* $\mathcal{G} = \langle L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_T, \Delta, \mathsf{wt}\rangle$ *where* $L = L_{\mathsf{Min}} \uplus L_{\mathsf{Max}} \uplus L_T$ *is a finite set of* locations*, partitioned into* Min *locations,* Max *locations, and* target *locations,* $\Delta \subseteq L \times \mathsf{Guards} \times \{\emptyset, \{x\}\} \times L$ *is a finite set of* transitions*, and* $\mathsf{wt}\colon \Delta \uplus L \to \mathbb{Z}$ *is a* weight function*.*

An example of WTG is depicted in Figure 1.

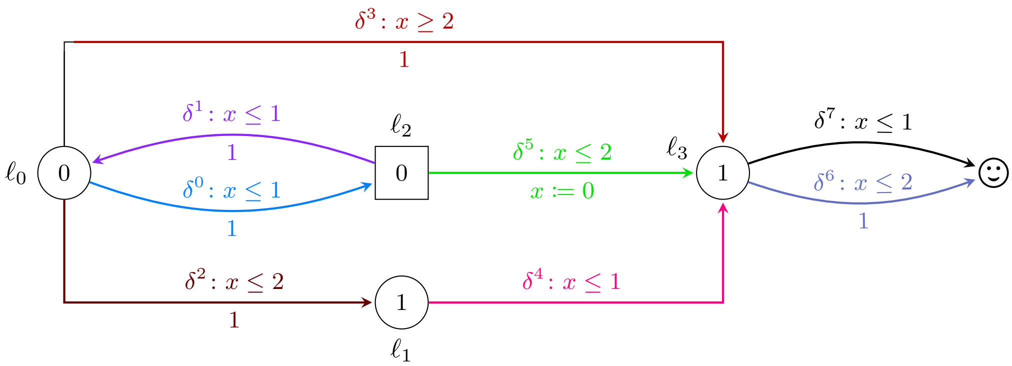

Figure 1 A WTG where circle locations are controlled by Min, the square location is controlled by Max and the target set is $T = \{☺\}$. Interval guards are denoted by constraints, for instance $x \geq 2$ represents $[2, +\infty]$. A reset of $Y = \{x\}$ is indicated by $x := 0$, and when $Y = \emptyset$ we ommit the reset on the transition. The weight function is indicated by labels on locations and transitions, *e.g.* $\mathsf{wt}(\ell_3) = 1$, $\mathsf{wt}(\delta^2) = 1$, and omitted weights on transitions are 0, *e.g.* $\mathsf{wt}(\delta^4) = 0$.

Before formalizing definitions, let us first give intuitions on how weighted timed games are played. We consider exact and robust semantics for weighted timed games, each given by a 2-player game played on an infinite transition system. In these games, each player chooses delays and transitions from configurations (i.e. a pair from a location and a valuation) associated with their respective locations. Of particular interest to us is the robust semantics, which intuitively, in contrast to the exact semantics, allows player Max to perturb the delays chosen by Min. More precisely, in the p-robust semantics, for a fixed parameter $p > 0$, Max will be allowed to perturb Min's delays by at most p. We adopt here the conservative notion of robustness (see *e.g.* [24]) in which the delays proposed by Min must remain feasible after Max has applied any possible perturbation. The exact semantics can be seen as the instantiation of the p-robust semantics with $p = 0$, and for $p > 0$, the p-robust conservative semantics may only reduce the choices of Min.

To define the exact and robust semantics, we introduce the following notations. *Configurations* of $\mathcal{G}$ are pairs $(\ell, \nu) \in L \times \mathbb{R}_{\geq 0}$ formed of a location and a valuation. $\mathsf{Conf} = \mathsf{Conf}_{\mathsf{Min}} \uplus \mathsf{Conf}_{\mathsf{Max}} \uplus \mathsf{Conf}_T$ denotes the set of all configurations, partitioned into configurations belonging to Min, Max and final configurations, according to the type of locations.

▶ **Definition 2.** *A WTG* $\mathcal{G} = \langle L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_T, \Delta, \mathsf{wt}\rangle$ *and* $p \geq 0$ *induce the* p-robust conservative semantics *(for* $p > 0$*) and* exact semantics *(for* $p = 0$*)* $[\![\mathcal{G}]\!]^p = \langle S, E^p, \mathsf{wt}\rangle$ *where*

- $S = S_{\mathsf{Min}} \uplus S_{\mathsf{Max}} \uplus S_T$ *is the set of* states *with* $S_{\mathsf{Min}} = \mathsf{Conf}_{\mathsf{Min}}$*,* $S_{\mathsf{Max}} = \mathsf{Conf}_{\mathsf{Max}} \cup (\mathsf{Conf}_{\mathsf{Min}} \times \Delta \times \mathbb{R}_{\geq 0})$ *and* $S_T = \mathsf{Conf}_T$*;*
- $E^p = E^p_{\mathsf{Min}} \uplus E^p_{\mathsf{Max}} \uplus E^p_{\mathsf{prtb}}$ *is the set of* edges *with*

$$E^p_{\mathsf{Max}} = \big\{\big((\ell,\nu) \xrightarrow{\delta,t} (\ell', (\nu+t)[Y := 0])\big) \mid \ell \in L_{\mathsf{Max}}, \delta = (\ell, I, Y, \ell') \in \Delta, \nu + t \models I\big\}$$
$$E^p_{\mathsf{Min}} = \big\{\big((\ell,\nu) \xrightarrow{\delta,t} ((\ell,\nu),\delta,t)\big) \mid \ell \in L_{\mathsf{Min}}, \delta = (\ell, I, Y, \ell') \in \Delta, [\nu+t, \nu+t+p] \subseteq I\big\}$$
$$E^p_{\mathsf{prtb}} = \big\{\big(((\ell,\nu),\delta,t) \xrightarrow{\delta,\varepsilon} (\ell', (\nu+t+\varepsilon)[Y := 0])\big) \mid \delta = (\ell, I, Y, \ell') \in \Delta, \varepsilon \in [0,p]\big\}$$

- $\mathsf{wt}\colon S \cup E^p \to \mathbb{Z}$ *the weight function such that for all states* $s \in S$ *with* $s = (\ell,\nu)$ *or* $s = ((\ell,\nu),\delta,t)$*,* $\mathsf{wt}(s) = \mathsf{wt}(\ell)$*, and all edges* $e \in E^p$*,* $\mathsf{wt}(e) = \mathsf{wt}(\delta)$ *if* $e = (s \xrightarrow{\delta,t} s')$ *with* $s \in \mathsf{Conf}$*, or* $\mathsf{wt}(e) = 0$*, otherwise.*

▶ Remark 3. In [7, 8, 16], the set of possible perturbations for Max forms the symmetric interval $[-p', p']$. For technical reasons, we use here an equivalent formalisation with an increase of at most $p = 2p'$ of the delays proposed by Min.

Note that states in S_{Max} are of two forms: either a configuration of Max in which they choose a delay and an action, or a state in $\mathsf{Conf}_{\mathsf{Min}} \times \Delta \times \mathbb{R}_{\geq 0}$, from which Max only chooses the delay perturbation via an edge in E^p_{prtb}. The latter will be referred to as *virtual states*, as they do not correspond to configurations of $\mathcal{G}$; all other states are *real states*.

In the special case of $p = 0$, the infinite transition system $[\![\mathcal{G}]\!]^0$ describes the exact semantics of the game. Perturbation edges in E^p_{prtb} are trivial: each step of Min with its choice in E^p_{Min} of delay and transition (δ, t), is followed by a dummy edge $((\ell,\nu),\delta,t) \xrightarrow{\delta,0} (\ell',\nu') \in E^p_{\mathsf{prtb}}$ corresponding to Max having no possibility to perturb the chosen delay t.

For s a state of $[\![\mathcal{G}]\!]^p$, we write $E^p(s)$ for the set of *enabled edges* in $[\![\mathcal{G}]\!]^p$ from s, that is, the edges with source state s. A state s is a *deadlock* when $E^p(s) = \emptyset$. We extend the notation E^p for the set of *enabled transitions* and terminology of deadlock to locations in $\mathcal{G}$.

Notably, the p-robust semantics may turn real states of Min into deadlocks even if the exact semantics contains an enabled edge. Thus, unlike in the literature [1], we allow state *and* location deadlocks.

Reaching as cheap as possible. In the game $\mathcal{G}$ under p-robust (or exact) semantics, the objective of Min is to reach a target configuration, while minimising the cumulated weight of the finite play built with Max.

A *finite play* of $\mathcal{G}$ w.r.t. the p-robust semantics is a sequence of edges in the transition system $[\![\mathcal{G}]\!]^p$ that starts in a real state of $\mathcal{G}$. Said differently, finite plays cannot start in a virtual state. We denote by $|\rho|$ the number of edges from $E^p_{\mathsf{Max}} \cup E^p_{\mathsf{Min}}$ of ρ (i.e. it is not exactly the length of the play), and by $\mathsf{last}(\rho)$ its last state A *maximal play* is a maximal sequence of consecutive edges: it is either a finite play reaching a deadlock state of $\mathcal{G}$, or an infinite sequence such that all its prefixes are finite plays. We use $\mathsf{FPlays}^p_{\mathcal{G}}$ to denote the set of finite plays on $[\![\mathcal{G}]\!]^p$, and write $\mathsf{FPlays}^p_{\mathcal{G},\mathsf{Min}}$ (resp. $\mathsf{FPlays}^p_{\mathcal{G},\mathsf{Max}}$) for the subset of finite plays ending in a state of Min (resp. Max).

To deal with the objective of Min, we associate to every finite play $\rho = s_0 \xrightarrow{\delta_0,t_0} s_1 \xrightarrow{\delta_1,t_1} \cdots s_k$ its *cumulated weight*, taking into account both discrete and continuous weights: $\mathsf{wt}_\Sigma(\rho) = \sum_{i=0}^{k-1}[t_i \times \mathsf{wt}(s_i) + \mathsf{wt}(\delta_i)]$. For a maximal play ρ, its *weight* $\mathsf{wt}(\rho)$ is set to $\mathsf{wt}_\Sigma(\rho)$ if it ends in a target real state (ℓ, ν) with $\ell \in L_T$, and $\mathsf{wt}(\rho) = +\infty$ if ρ does not reach L_T (either ρ is infinite or it ends in a deadlock state that is not a target state).

A *path* is a finite or infinite sequence π of transitions of $\mathcal{G}$ and we write $\mathsf{FPaths}_{\mathcal{G}}$ for the set of finite paths. A *target path* is a finite path ending in the target set L_T. We denote by $\pi_1 \xrightarrow{\delta} \pi_2$ the concatenation of two paths $\pi_1 \xrightarrow{\delta}$ and π_2, and by $|\pi|$ the length of π. Note that each play ρ of $[\![\mathcal{G}]\!]^p$ projects to a unique path π (by keeping only the transitions of edges from real states of $\mathcal{G}$): we say that ρ *follows* path π.

A *strategy* for Min (resp. Max) is a mapping from finite plays ending in a state of Min (resp. Max) to a decision in (δ, t) labelling an edge of $[\![\mathcal{G}]\!]^p$ from the last state of the play. As argued earlier, states of Min may contain deadlock states, so that we consider strategies of Min to be partial mappings. For instance, in the WTG depicted in Figure 1 and a perturbation p, a strategy for Min in all plays ending in (ℓ_3, ν) can be defined only when $\nu(x) \leq 2 - p$ since, otherwise, there are no outgoing edges in $[\![\mathcal{G}]\!]^p$ from this state. Symmetrically, we require Max to always propose a move unless the play is in a deadlock state - since otherwise Max can create deadlock by not playing. More formally, a strategy for Min, denoted σ^p, is a (possibly partial) mapping $\sigma^p \colon \mathsf{FPlays}^p_{\mathsf{Min}} \to E^p$ such that $\sigma^p(\rho) \in E^p(\mathsf{last}(\rho))$. A strategy for Max, denoted τ^p, is a (possibly partial) mapping $\tau^p \colon \mathsf{FPlays}^p_{\mathsf{Max}} \to E^p$ such that for all ρ, if $E^p(\mathsf{last}(\rho)) \neq \emptyset$, then $\sigma^p(\rho)$ is defined, and in this case $\sigma^p(\rho) \in E^p(\mathsf{last}(\rho))$. The set of strategies of Min (resp. Max) with the perturbation p is denoted by $\mathsf{Strat}^p_{\mathcal{G},\mathsf{Min}}$ (resp. $\mathsf{Strat}^p_{\mathcal{G},\mathsf{Max}}$). If the game is clear from the context, we remove $\mathcal{G}$ from (all) the notations.

A play or finite play $\rho = s_0 \xrightarrow{\delta_0,t_0} s_1 \xrightarrow{\delta_1,t_1} \cdots$ *conforms* to a strategy σ^p of Min if for all k such that s_k belongs to Min, we have that $(\delta_k, t_k) = \sigma^p(s_0 \xrightarrow{\delta_0,t_0} \cdots s_k)$. Similarly, we define when plays conform to strategies τ^p of Max. For every pair of strategies (σ^p, τ^p) of players Min and Max, and for every real state (ℓ_0, ν_0), we let $\mathsf{Play}((\ell_0, \nu_0), \sigma^p, \tau^p)$ be the *outcome* of σ^p and τ^p from (ℓ_0, ν_0), defined as the unique maximal play that starts in (ℓ_0, ν_0) and conforms to σ^p and τ^p.

Robust value problem. We are interested in the minimal cumulated weight that Min can guarantee while reaching the target, whatever the choices of actions and perturbations Max makes. To formalize this, we introduce *robust values*. For $p > 0$, (ℓ, ν) a real state of $\mathcal{G}$ and $\tau^p \in \mathsf{Strat}^p_{\mathsf{Max}}$ a fixed strategy for Max, we let $\mathsf{rVal}^{p,\tau^p}_{\mathcal{G}}(\ell, \nu) = \inf_{\sigma^p \in \mathsf{Strat}^p_{\mathsf{Min}}} \mathsf{wt}(\mathsf{Play}((\ell, \nu), \sigma^p, \tau^p))$ be the minimal cumulated weight Min can ensure assuming Max plays τ^p. Then the *p-robust* Max *value* is $\underline{\mathsf{rVal}}^p_{\mathcal{G}}(\ell, \nu) = \sup_{\tau^p \in \mathsf{Strat}^p_{\mathsf{Max}}} \mathsf{rVal}^{p,\tau^p}_{\mathcal{G}}(\ell, \nu)$. (Note however that this value is not defined for virtual states of the robust semantics.) Since the p-robust conservative semantics defines a *quantitative reachability game* [9], the general determinacy result of [10, Theorem 2.2] applies and the *p-robust value* can be defined as follows: $\mathsf{rVal}^p_{\mathcal{G}}(\ell, \nu) = \underline{\mathsf{rVal}}^p_{\mathcal{G}}(\ell, \nu) = \inf_{\sigma^p \in \mathsf{Strat}^p_{\mathsf{Min}}} \sup_{\tau^p \in \mathsf{Strat}^p_{\mathsf{Max}}} \mathsf{wt}(\mathsf{Play}((\ell, \nu), \sigma^p, \tau^p))$.

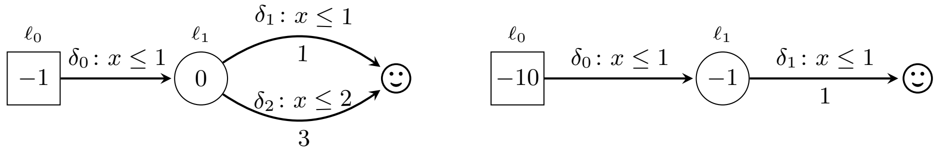

Figure 2 Left: a WTG where Max can increase the value by perturbing the delays of Min. Right: a WTG in which the exact and robust values coincide.

For $p = 0$, rVal^0 coincides with the *(exact) value* at the core of the well-studied value problem for WTGs [6, 5, 23, 17, 18]. We sometimes write Val for rVal^0 for simplicity. Moreover, the values enjoy the following monotony property:

▶ **Lemma 4** ([22]). *Let $\mathcal{G}$ be a WTG, and $p > p' \geq 0$ be two perturbations. Then, for every real state (ℓ, ν), $\mathsf{rVal}^p(\ell, \nu) \geq \mathsf{rVal}^{p'}(\ell, \nu)$.*

In this paper, we are interested in the value that Min can guarantee if Max is restricted to arbitrarily small perturbations. Rather than a fixed perturbation $p > 0$, the *robust value problem* is defined for infinitesimal perturbations. Given a WTG $\mathcal{G}$, the *robust value* is defined, for every real state (ℓ, ν) of $\mathcal{G}$, by $\mathsf{rVal}^{0^+}(\ell, \nu) = \lim_{p \to 0, p>0} \mathsf{rVal}^p(\ell, \nu)$. The robust value function is defined as the limit of the p-robust value functions, and it is well-defined since it is the limit of a non-increasing sequence of functions (see Lemma 4).

▶ **Example 5.** To illustrate the notions of values, consider the WTG depicted in Figure 2 (left). From the initial real state $(\ell_0, 0)$, this is a typical example where the exact and robust values do not coincide. For the exact value, we observe that Max has no interest in staying in location ℓ_0. Indeed, whatever the delay chosen by Max in ℓ_0, Min will always choose δ_1 from ℓ_1, and waiting in ℓ_0 decreases the accumulated weight, which is profitable to Min. Thus $\mathsf{Val}(\ell_0, 0) = 1$. For the robust value however, intuitively Max can force Min to take δ_2 by reaching ℓ_1 with clock value 1. If we fix $p > 0$, from $(\ell_1, 1)$, Min cannot choose δ_1 even with 0 delay, as $[1, 1 + p]$ is not included in where the guard of δ_1. Spending 1 time unit in ℓ_0, Max ensures $\mathsf{rVal}^p(\ell_0, 0) \geq 2$. In fact $\mathsf{rVal}^p(\ell_0, 0) = 2 + p$ and thus $\mathsf{rVal}^{0+}(\ell_0, 0) = 2$.

For the WTG in Figure 2 (right), the gap between the exact and robust values is even infinite! The exact value from $(\ell_0, 0)$ corresponds to the behaviour where Max leaves ℓ_0 immediately, and Min stays 1 time unit in ℓ_1. Thus $\mathsf{Val}(\ell_0, 0) = -1$. Now, fixing a perturbation $p > 0$, it is in Max's interest to reach ℓ_1 with clock value $1 - p + \varepsilon$, for any $\varepsilon \in]0, p]$. Doing so, Max blocks Min in ℓ_1: indeed, Min cannot propose any delay, as even for delay 0, the interval $[1 - p + \varepsilon, 1 + \varepsilon]$ is not included in the guard of the unique outgoing transition δ_1. Therefore, Max can prevent Min to reach the target and $\mathsf{rVal}^{0^+}(\ell_0, 0) = +\infty$.

The decision problem associated with the robust value is called the *robust value problem.* It is naturally defined as follows: given a WTG $\mathcal{G}$, an initial location ℓ_0, and a threshold $\lambda \in \mathbb{Q}$, does $\mathsf{rVal}^{0^+}(\ell_0, 0) \leq \lambda$?

Unsurprisingly, this problem is undecidable when generalized to WTGs with multiple clocks [7, Theorem 4]. Some recent results [22] prove the decidability for restricted classes of WTGs, for instance WTGs that are acyclic [1] or divergent [6, 12]. We focus here on another class of WTGs for which the exact value problem is known to be decidable [23], namely one-clock WTGs. As in [22], we consider WTGs in which no state has (exact) value $-\infty$, and such that the clock valuation is bounded by a maximal constant $\mathsf{M} \in \mathbb{N}$. Under these two common assumptions, our main contribution is the following:

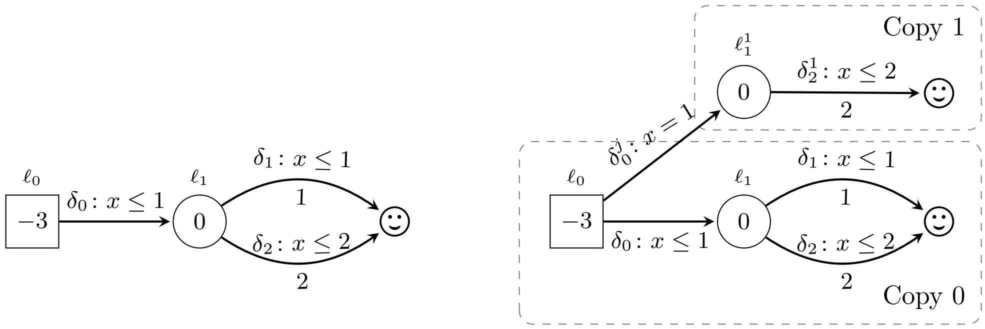

Figure 3 A WTG (left) where Max has no incentive to block Min to play optimally and the corresponding copy game (right). Max has a new jumping transition δ_o^j from ℓ_0 to ℓ_1^1 they can use to prevent Min to play δ_1 in the next move: δ_1 is no longer present in copy 1.

▶ **Theorem 6.** *The robust value problem is decidable for one-clock WTGs.*

The high-level proof idea to establish this decidability result is as follows.

1. Since we saw in Example 5 that robustness may introduce deadlocks, our first step is to discriminate paths[1] that are feasible or unfeasible in the p-robust semantics. Typically, the path $\ell_0 \xrightarrow{\delta_0} \ell_1 \xrightarrow{\delta_1}$ on the WTG right of Figure 2 is p-unfeasible, for any $p > 0$.
2. Building on the characterization of r-feasible paths (paths that are p-feasible for small values of p), from a game $\mathcal{G}$, we construct another 1-clock WTG $\mathcal{G}'$ with same robust value, and in which every path is either r-feasible or unfeasible (that is, 0-unfeasible). A naive construction of $\mathcal{G}'$ that would keep only r-feasible paths is however not correct. To illustrate this, consider the WTG left of Figure 3. This example only differs from the one left of Figure 2 in the weights in ℓ_0 and on δ_2. However, because of the smaller weight -3 in ℓ_0, it is not profitable to Max to block δ_1, and one can show that $\mathsf{rVal}^{0^+}(\ell_0, 0) = \mathsf{Val}(\ell_0, 0) = 1$. Therefore, simply removing δ_1 would not preserve the robust value: it would increase to $\mathsf{rVal}^{0^+}(\ell_0, 0) = 2$. Instead, in our construction, $\mathcal{G}'$ consists in several copies of $\mathcal{G}$ indexed by integer clock values (up to the maximal constant) and Max can choose to jump to a copy with larger index, say $i \in \mathbb{N}$, to mimick the fact that they want to use robustness and block transitions with guards I such that $\mathsf{right}(I) \leq i$. We call it the *copy game*. Figure 3 illustrates the construction on the above-mentioned WTG. From $(\ell_0, 0)$ Max may choose to jump to copy 1 to prevent Min to play δ_1, but they may as well stay in copy 0.
3. Finally, in the copy game, we prove that the robust value coincides with the exact value. We therefore provided a polynomial reduction from the robust value problem in $\mathcal{G}$ to the exact value problem in $\mathcal{G}'$, which is decidable.

3 Characterization of Robust Deadlocks

The first step of the proof of Theorem 6 consists in the identification of paths in the p-robust semantics where Max can exploit delay perturbation to create a deadlock, as illustrated in Example 5. They do so by choosing a perturbed delay such that the reached valuation is close enough to the upper-bound of the guard of Min's next transition. Our objective is to determine when these deadlocks due to perturbations are profitable for Max. We do so directly on the WTG structure, i.e. only by inspecting its paths.

[1] Recall that paths are sequence of transitions in the game description $\mathcal{G}$.

In the following, we fix a WTG $\mathcal{G}$. We assume that $\mathcal{G}$ does not contain transitions with a singleton guard from locations of Min. This is without lost of generality, since Min cannot choose a transition with a singleton guard otherwise any positive delay perturbation of Max would enforce the guard to be unsatisfied.

We introduce the notions of *feasible* and *r-feasible* paths (resp. unfeasible and r-unfeasible paths) representing the fact that Max cannot constrain, via any choice of delay, Min to be unable to take the next transition on the path. Said differently, a path is feasible when Max cannot avoid all plays along this path. To properly define this notion, we define the *sub-game induced by a path* π by the unfolded WTG, denoted by $\mathcal{G}_\pi$, where locations are exactly the ones of π (locations from $\mathcal{G}$ may have many occurrences in $\mathcal{G}_\pi$), and its transitions are only ones of π.

▶ **Definition 7.** *A path π of $\mathcal{G}$ is* r-unfeasible *if there exists $p \geq 0$ and $\tau^p \in \mathsf{Strat}^p_{\mathcal{G}_\pi,\mathsf{Max}}$ such that for every strategy $\sigma^p \in \mathsf{Strat}^p_{\mathcal{G}_\pi,\mathsf{Min}}$ of* Min *in $\mathcal{G}_\pi$, we have*[2]*:* $|\mathsf{rPlays}^p((\ell_0, 0), \sigma^p, \tau^p)| \neq |\pi|$ *where, ℓ_0 is the first location of π. When $p = 0$, we simply say that π is* unfeasible.

For instance, the path $\ell_0 \xrightarrow{\delta_0} \ell_1 \xrightarrow{\delta_1} ☺$ of the WTG depicted in Figure 2 (right) is feasible since Min can always reach the target. However, it is r-unfeasible (since for every $p > 0$, the strategy of Max reaching $(\ell_1, 1)$ in the first step avoids Min to take δ_1 as $E^p(\ell_1, 1) = \emptyset$.

Although the r-feasibility of a path quantifies over all strategies of Min, fortunately, we observe that an r-unfeasible path is blocking already in the extremal case when Min's strategy is to play minimum delays. Thus, checking whether a path is r-unfeasible can be done on such a strategy. Note however that the minimum delay may not always exist (if the guard interval is left-open), thus we define a family of strategies for Min indexed by smaller and smaller errors ε. Formally, given a path $\pi = \ell_0 \xrightarrow{\delta_0} \dots \xrightarrow{\delta_{k-1}} \ell_k$, a perturbation bound p, and $\varepsilon > 0$, we define an *asap strategy* of Min on $\mathcal{G}_\pi$, denoted by $\sigma^{p,\varepsilon}_\pi$, such that for all plays $\rho \in \mathsf{rPlays}^p_{\mathsf{Min}}$ where ρ follows π and $|\rho| < |\pi|$:

$$\sigma^{p,\varepsilon}_\pi(\rho) = \begin{cases} (\delta_{|\rho|}, t_{\min}(\nu_{|\rho|-1}, I_{|\rho|})) & \text{if } \nu_{|\rho|-1} + t_{\min} \models I_{|\rho|}; \\ (\delta_{|\rho|}, t_{\min}(\nu_{|\rho|-1}, I_{|\rho|}) + \varepsilon) & \text{otherwise}; \end{cases}$$

where $I_{|\rho|}$ is the guard of the transition $\delta_{|\rho|}$, and, for all valuations ν and guards I, we define $t_{\min}(\nu, I) = \inf\{t \mid \nu + t \in \overline{I}\}$. As the intuition suggests, one can show that asap strategies suffice to check r-unfeasibility of paths.

▶ **Lemma 8.** *Let $(\sigma^{p,\varepsilon}_\pi)_\varepsilon$ be a family of asap strategies for* Min *for a path π starting at location ℓ_0. Then, the two following statements are equivalent:*

1. *π is r-unfeasible;*
2. *there exists $p > 0$, a robust strategy τ^p for* Max *in $\mathcal{G}_\pi$ and $\epsilon > 0$ such that for all $0 < \widehat{\varepsilon} \leq \varepsilon$,* $|\mathsf{rPlay}^p((\ell_0, 0), \sigma^{p,\widehat{\varepsilon}}_\pi, \tau^p)| < |\pi|$.

Since unfeasible paths can be detected by a reachability query in the induced timed game [20], we now have all ingredients to characterize paths that are feasible but r-unfeasible. Indeed, deadlocks induced by robustness are of two types:

type 1 Min cannot guarantee a sufficient margin with the upper-bound of the guard I on one of its transitions, because since the last reset, a guard I' on one of its transition imposes a valuation greater or equal than the upper-bound of I;

[2] We recall that, given a play ρ, $|\rho|$ denotes its number of edges except the ones of E^p_{prtb}, that is, the length of its underlying path.

type 2 Max can exploit delays by pushing the clock valuation towards the upper-bound of the guard I of one of its transitions, and there are no reset before a transition of Min whose guard I' has a smaller upper-bound than I.

We now formalize these two possible patterns for feasible but r-unfeasible paths, using $\mathsf{Paths}_{\mathsf{Min}}^{Y=\emptyset}$ to denote the set of paths in $\mathcal{G}$ without reset transitions ending in a state of Min.

▶ **Proposition 9.** *Let π be a feasible path in $\mathcal{G}$. The path π is r-unfeasible if and only if it has the following form: $\pi = \pi_0 \xrightarrow{\delta_0} \pi_1 \xrightarrow{\delta_1} \pi_2$ with $\pi_1 \in \mathsf{Paths}_{\mathsf{Min}}^{Y=\emptyset}$, and*

type 1 $\pi_0 \in \mathsf{Paths}_{\mathsf{Min}}$ *and* $\mathit{left}(I_0) = \mathit{right}(I_1)$*; or*

type 2 $\pi_0 \in \mathsf{Paths}_{\mathsf{Max}}$ *and* $\mathit{right}(I_0) = \mathit{right}(I_1)$*;*

with I_i the guard of δ_i for $i \in \{0, 1\}$.

4 Transformation into the Copy Game

The second step in the proof of Theorem 6 is the construction from $\mathcal{G}$ of a new WTG, the *copy game*, whose exact value coincides with the robust value in $\mathcal{G}$. In this WTG, Max can either choose to play an r-unfeasible path or stick to a feasible path. As explain in Section 2, it would be incorrect to remove r-unfeasible paths (or parts of such paths) without giving extra power to Max. Specifically, the copy game allows Max to emulate a robust deadlock in the exact semantics (without impacting them when they prefer not to use robust deadlocks). As an illustrative example, the WTG depicted in Figure 4 is the associated copy game of the WTG depicted in Figure 1.

Before coming to the formal definition, let us explained intuitively how the construction proceeds. The copy game is composed of different copies of the WTG, one for each *open region*[3]. Each original transition of Min with guard I is duplicated in the copy a (for the interval $(a, +\infty)$) as a transition with guard $I \cap [a, +\infty)$ and with destination the highest copy between a and the lower-bound of I (only when a is smaller than the upper-bound of I). Each original transition of Max is transformed in copy a to a transition with guard $I \cap [a, +\infty)$, and also to any *jumping transition*, denoted with the superscript j, with destination in copy b (with $b \neq a$) with guard $\{b\}$. Formally, an original transition $\delta = (\ell, I, Y, \ell') \in \Delta$ can be transformed into two different ways:

$$\delta_{a\to b} = (\langle \ell, a\rangle, I \cap [a + \infty), Y, \langle \ell', b\rangle) \qquad \delta^j_{a\to b} = (\langle \ell, a\rangle, \{b\}, Y, \langle \ell', b\rangle)$$

In the copy game, the two players have different abilities regarding the moves between copies. On the one hand, Min cannot choose the next copy: they are constrained by the minimal open region they must use. On the other hand, for non-resetting transitions Max may choose which copy to move to: they may use a jumping transition $\delta^j_{a\to b}$ for any b in the guard closure.

We now formally define the copy game.

▶ **Definition 10.** *For $\mathcal{G}$ a WTG with maximal constant* M*, its associated* copy game *is $\mathcal{C} = \langle L^{\mathcal{C}}, \Delta^{\mathcal{C}}, \mathsf{wt}^{\mathcal{C}}\rangle$ defined by*

- $L^{\mathcal{C}} = L \times [\![0, M]\!]$ *;*

[3] In single-clock WTGs, the classical notion of regions reduces to simple intervals on $\mathbb{R}$, see [6] for details. Here, we call open region, the infinite interval $(a, +\infty)$ where a is a lower-bound of a region.

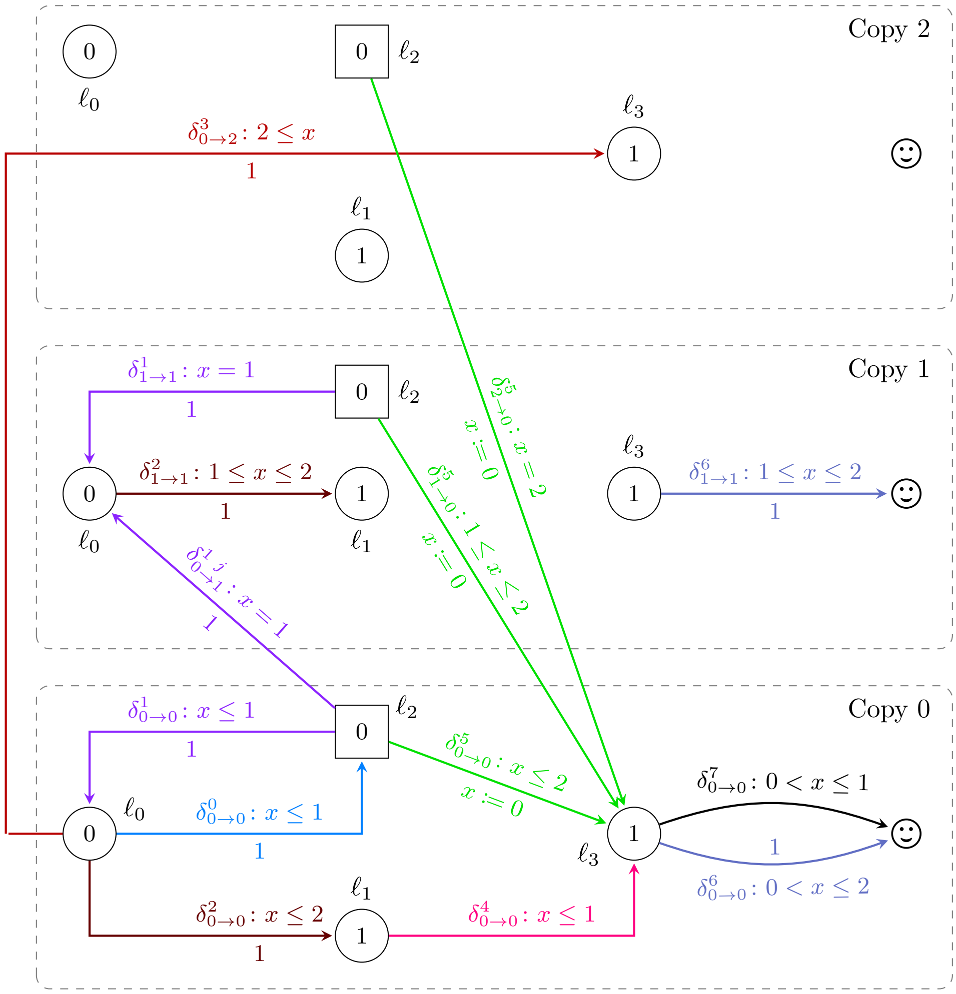

Figure 4 The copy game associated with the WTG depicted in Figure 1.

- $\Delta^{\mathcal{C}} = \bigcup_{a\in[\![0,M]\!]} \cup_{\delta\in\Delta} \Delta^{\mathcal{C}}_a(\delta)$ *where, for all* $a \in [\![0, M]\!]$, *and all* $\delta = (\ell, I, Y, \ell')$, *we fix*

$$\Delta^{\mathcal{C}}_a(\delta) = \begin{cases} \{\delta_{a\to\max(a,\mathsf{left}(I))}\} & \text{if } \ell \in L_{\mathsf{Min}},\ Y = \emptyset, \text{ and } \mathsf{right}(I) > a \\ \{\delta_{a\to 0}\} & \text{if } \ell \in L_{\mathsf{Min}},\ Y = \{x\}, \text{ and } \mathsf{right}(I) > a \\ \{\delta_{a\to a}\} \cup \{\delta^j_{a\to b} \mid b \in \overline{I} \cap [\![a+1, M]\!]\} & \text{if } \ell \in L_{\mathsf{Max}},\ Y = \emptyset, \text{ and } \mathsf{right}(I) \ge a \\ \{\delta_{a\to 0}\} & \text{if } \ell \in L_{\mathsf{Max}},\ Y = \{x\}, \text{ and } \mathsf{right}(I) \ge a\ ; \end{cases}$$

- $\mathsf{wt}^{\mathcal{C}}(\langle \ell, a\rangle) = \mathsf{wt}(\ell)$, *for all* $\langle \ell, a\rangle \in L^{C}$, *and* $\mathsf{wt}^{\mathcal{C}}(\delta_{a\to b}) = \mathsf{wt}^{\mathcal{C}}(\delta^j_{a\to b}) = \mathsf{wt}(\delta)$, *for all transitions of* $\Delta^{\mathcal{C}}$.

Note that the constraint $\mathsf{right}(I) > a$ in the transitions from Min locations ensures that the guard is neither empty and nor reduced to a singleton. This is convenient since such transitions would not be possible for Min in the robust semantics. Consequently, the copy game exhibits no robust deadlocks of type 1. Moreover, the number of copies is exponential in the size of $\mathcal{G}$ (since M is encoded in binary). Finally, the copy index represents the infimum valuation of plays because the lower-bound of transitions guards between two copies is the index of the highest copy, and resetting transitions have destination copy 0. These interesting properties are summarized in the following lemma.

▶ **Lemma 11.** *Let $\mathcal{C}$ be the copy game associated with $\mathcal{G}$. Then,*
1. *there are no feasible and r-unfeasible paths of type 1 in $\mathcal{C}$;*
2. *$\mathcal{C}$ is exponential in the size of $\mathcal{G}$;*
3. *for every finite play starting in the 0-th copy and ending in a location $\langle \ell, a\rangle$, its last valuation ν satisfies $\nu \geq a$.*

5 Recovering the Original Robust Value in the Copy Game

To conclude the proof of Theorem 6, we show the correctness of the reduction to the copy game. More precisely, the robust value in $\mathcal{G}$ is equal to the (exact) value in $\mathcal{C}$. This not only holds for initial states of the form $(\ell_0, 0)$ but also for any (ℓ, ν), provided ν is either 0 or not integer. Formally:

▶ **Theorem 12.** *Let (ℓ, ν) be a real state such that $\nu \notin \mathbb{N}^*$, $\mathsf{rVal}^{0+}_{\mathcal{G}}(\ell, \nu) = \mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu)$.*

Before proceeding to the proof of Theorem 12, let us justify how it allows one to conclude the proof of Theorem 6. Let $\mathcal{G}$ be a WTG and $\mathcal{C}$ be its associated copy game. Thanks to Theorem 12, solving the robust value problem in $\mathcal{G}$ for $(\ell_0, 0)$ reduces to solving the exact value problem in $\mathcal{C}$ from $(\langle \ell_0, 0\rangle, 0)$ in copy 0. Since computing the exact value in a WTG with one-clock can be done in exponential time [23], we derive that the robust value problem in $\mathcal{G}$ can be solved in exponential time in the size of $\mathcal{C}$. Now, $\mathcal{C}$ is exponential in the size of $\mathcal{G}$ (by Lemma 11), so that we conclude that the robust value problem is solvable in doubly exponential time in the size of $\mathcal{G}$. In particular, the robust value problem is decidable for 1-clock WTG. This concludes the proof of Theorem 6.

The rest of this section presents the key ingredients to prove Theorem 12. The proof is in two steps. First, we prove that the construction of the copy game preserves the *robust* value. We do so by showing the existence of a perturbation bound such that the fixed-perturbation robust values in the two games are arbitrarily close. Second, we prove that in the copy game, the robust value is equal to the (exact) value, as far as one considers for initial valuation either 0, or any valuation in a left-open interval.

Preservation of the robust value

The first step of the correctness proof of the copy game consists in proving the preservation of the robust value between the two WTGs $\mathcal{G}$ and $\mathcal{C}$. More precisely, we prove that from a given robust strategy of Max for one of these games, one can define an equivalent one (i.e. with an at least as good value) in the other game. Both transformations are technical yet rather natural.

▶ **Proposition 13.** *Let (ℓ, ν) be a real state. Then, $\mathsf{rVal}^{0+}_{\mathcal{G}}(\ell, \nu) = \mathsf{rVal}^{0+}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu)$.*

The first inequality $\mathsf{rVal}^{0+}_{\mathcal{G}}(\ell, \nu) \leq \mathsf{rVal}^{0+}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu)$ is the most natural, since $\mathcal{C}$ does not constrain the robust behaviours of Max. In particular, we can naturally embed a robust strategy of Max in $\mathcal{G}$ into one of $\mathcal{C}$ by ignoring all jumping transitions. The plays conforming to this new strategy in $\mathcal{C}$ will remain in copy 0.

The second inequality $\mathsf{rVal}^{0+}_{\mathcal{G}}(\ell, \nu) \geq \mathsf{rVal}^{0+}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu)$ is more involved, since Max may use jumping transitions in $\mathcal{C}$ with clock valuation that does not satisfy the guard of the original transition in $\mathcal{G}$. Indeed, by definition of $\mathcal{C}$, if I is the guard in $\mathcal{G}$, then the guard $\{b\}$ of the jumping transition only ensures $b \in \overline{I}$, yet we may have $b \notin I$. In particular, when Max chooses such a transition in $\mathcal{C}$, to mimick it in $\mathcal{G}$, one needs to slightly alter the delay (so that it satisfies the guard I). However, since the robust value rVal^p is, a priori, not continuous, such a deviation must be compensated in the rest of the play. Ultimately, the plays in $\mathcal{C}$ and $\mathcal{G}$ should reacht the exact same valuation. To control precisely the potential weight difference (due to a delay difference), this delay alteration must depend on the largest location weight in $\mathcal{G}$, and also on the length of the play.

Interestingly, the proof of the first inequality further induces that Max can play optimally without jumping transitions in $\mathcal{C}$, i.e.:

▶ **Corollary 14.** *Let $p > 0$ and $\varepsilon > 0$. Then,* Max *has an ε-optimal robust strategy in* $\mathsf{Strat}^p_{\mathcal{C},\mathsf{Max}}$ *such that all its conforming plays contain no jumping transitions.*

Exact and robust values coincide in the copy game

The next step is to prove that the robust and the exact values are equal in $\mathcal{C}$.

▶ **Proposition 15.** *Let $\langle \ell, 0\rangle$ be a location of the 0-th copy of $\mathcal{C}$. Then, for every $\nu \notin \mathbb{N}^*$,* $\mathsf{rVal}^{0+}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) = \mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu)$.

This result constitutes the most technical part of this paper and forms the core of our contribution. The high-level intuition is that the robust and exact values coincide because the copy game $\mathcal{C}$ does not contain any path that is feasible and r-unfeasible. Compared to similar decidability results for the robust value problems, an additional layer of difficulty arises in our case with the copy game, and establishing the equality on the copy game is non trivial. One needs to massage strategies, to prove that optimal strategies for each of the players can be taken in a restricted class of *well-formed* strategies. Roughly speaking, Max can encode deadlocks that are incurred by robustness into classical deadlocks, by using jumping transitions instead of playing very close to the upper bound of a guard. For Min, restricting to well-formed strategies is less immediate and relies on the continuity of the classical value. Thanks to continuity, Min can avoid taking the risk of playing close to the bound of guards without significantly altering the weight; and since the plays are finite, Min can control the difference in weight between robust and exact plays. The rest of this section presents in more details the successive steps of the proof of Proposition 15.

To prove the proposition, we show the following result: there exists an integer $N \in \mathbb{N}$ such that for every $\varepsilon > 0$ and every location $\ell \in L$, there exists a perturbation bound $p > 0$ such that $\forall i \in \mathbb{N}^\star$ if $\nu \in]i, i+1-N\cdot 2p] \cup \{0\}$, then

$$\mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) \leq \mathsf{rVal}^p_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) \leq \mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) + \varepsilon \ . \tag{1}$$

Taking the limit when p tends to 0, we obtain, for all real states $(\langle \ell, 0\rangle, \nu)$ such that $\nu \notin \mathbb{N}^\star$, that $\mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) \leq \mathsf{rVal}^{0+}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) \leq \mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) + \varepsilon$. Since this inequality hold for all ε, we conclude that $\mathsf{rVal}^{0+}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) \leq \mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu)$.

Now, we successively prove the two inequalities from Equation (1). First, by [23], we already have that for all $p > 0$, and $(\langle \ell, 0\rangle, \nu)$ real states of $\mathcal{C}$, $\mathsf{rVal}^p_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) \geq \mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu)$. Thus, $\mathsf{rVal}^{0+}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) \geq \mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu)$.

Conversely, we prove the following statement. For $\varepsilon > 0$, there exist an integer $N \in \mathbb{N}$ and a perturbation bound $\eta > 0$ such that for all $p \leq \eta$, and real states $(\langle \ell, 0\rangle, \nu)$, such that $\nu \in (i, i+1-N\cdot p] \cup \{0\}$ with $i \in \mathbb{N}$, we have $\mathsf{rVal}^p_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) \leq \mathsf{Val}_{\mathcal{C}}(\langle \ell, 0\rangle, \nu) + \varepsilon$. To do so, we first fix N and η by considering the maximal length of plays to guarantee that Min can achieve the value. Then, we prove in Lemma 16 (resp. Lemma 17) that Max (resp. Min) can play almost optimally far enough of positive integer valuations in $\mathcal{C}$ (resp. $\mathcal{G}$) by enforcing jumping transitions where a valuation is too close to a guard upper-bound (resp. by the continuity of the exact value [10]). Finally, in Lemmas 18 and 19, we embed an exact strategy of Max into a robust one ensuring a value at least as large.

Determining the constants N and η. Intuitively, N is related to the valuations margin required by Min to guarantee that the target remains reachable under a robust semantics, independent of the perturbation chosen by Max. A result from robust reachability in timed games [8, Proposition 3.2], implies that there exists $N \in \mathbb{N}$ such that for every reachable pair of location and open region, the region contains an interval $(i, i+1-N \cdot p[$ when $p \leq \frac{1}{3N}$. In more details, [8] exploits the data structure of *shrunk DBM* [26] that builds upon Difference Bound Matrices (DBM), and which consists in a DBM constrained with a shrinking matrix indicating how far Min has to remains from border of regions. N is set as by the maximal element of this matrix and is computable in EXPTIME [8][4].

Now, the constant η must account for the weight alteration due to the perturbations of Max. To define η, we use the maximal bound on length of plays, denoted κ. Following [23], under the exact semantics, there exists κ such that for every $\epsilon > 0$, there exists an ϵ-optimal strategy σ for Min, for which the length of all conforming plays is bounded by κ. Formally, it is such that for every $\tau \in \mathsf{Strat}^0{}_{\mathcal{C},\mathsf{Max}}$, for every real state $(\langle \ell, 0\rangle, \nu)$, $|\mathsf{Play}((\langle \ell, 0\rangle, \nu), \sigma, \tau)| \leq \kappa$. We observe that κ also bounds the length of plays such that Min reaches the target with an optimal weight both under the exact and robust semantics, because the (exact) value is always at most any robust value (see Lemma 4). We let $\eta < \dfrac{\varepsilon}{24N^2 \cdot W_{\mathsf{loc}} \cdot \kappa}$ where W_{loc} is the maximal weight in locations of $\mathcal{C}$. By definition, $\eta < \frac{1}{3N}$ where the latter in the bound derived from the (unweighted) reachability objective that defines N.

Finally, we fix $p \leq \eta$ and $(\langle \ell, 0\rangle, \nu)$ be a real state of $\mathcal{C}$ such that $\nu \in (i, i+1-N \cdot p] \cup \{0\}$.

Restricting to relevant strategies of Max. The copy game $\mathcal{C}$ introduces *jumping transitions* that give Max the opportunity to choose a singleton guard to reach a copy with higher index, in which the options of Min can be limited compared to the original game. We will show that in $\mathcal{C}$, it is in the interest of Max–or rather they lose close to nothing– to choose a jumping transition instead of staying in the same copy but close to an integer valuation. In other words, Max can play almost optimally with a so-called *well-formed* strategy, guaranteeing that if they stay in the same copy, then the valuation remains bounded away from integers.

Formally, a robust strategy of Max, τ^p is *p-well-formed* if for every robust play ρ ending in a real state of Max, denoted $(\langle \ell, a\rangle, \nu)$, such that $\tau^p(\rho) = (\delta_{a \to a}, t)$, then, for every $i \in \mathbb{N}^*$, $\nu + t \notin F_i$ where we define $F_i = (i - N \cdot p, i)$ as the *forbidden interval* for $i \in \mathbb{N}$. Max can play optimally by picking strategies within the class of well-formed strategies, by forcing a jumping transition when the original valuation would fall in a forbidden interval. From any almost-optimal robust strategy of Max that does not take any jumping transition (by Corollary 14), we define a new strategy that will take a jumping transition when a forbidden interval is reached. The choice of the size of the interval will control the weight alteration incurred by jumping to a higher copy, so that the value this new strategy induces is close enough to the one of the original strategy. Thus, we obtain the following lemma.

▶ **Lemma 16.** Max *has a p-well-formed ε-optimal strategy in $\mathcal{C}$.*

In the sequel, we therefore consider only p-well-formed robust strategy for Max in $\mathcal{C}$.

[4] Reference [8] considers a robustness notion, called *excessive semantics*, that slightly differs from ours, in that the guard is checked before the perturbation. However, a polynomial reduction in the two directions between both semantics was later provided [22]. Proposition 3.2 of [8] thus applies to our conservative semantics as well.

Restricting to relevant strategies of Min. Against a p-well-formed strategy of Max, Min always has a robust decision along a feasible path in $\mathcal{C}$ (which is r-feasible by Lemma 11-Item 1). Thus, a crux of the proof is to show that we can restrict the choices of Min to robust ones in $\mathcal{C}$ without impacting too much the value. To control the value alteration, we use the fact that the (exact) value is continuous over regions [10].

Formally, an (exact) play is *p-well-formed* if for every edge $(\ell,\nu) \xrightarrow{(\delta,t)} (\ell',\nu')$ along the play, we have $t=0$ or $\nu+t \notin F_i$ for all $i \in \mathbb{N}^*$ (F_i is again the forbidden set for interval $(i,i+1)$). Similar to Max in the previous paragraph, one can prove that Min has an almost-optimal p-well-formed strategy in $\mathcal{C}$.

▶ **Lemma 17.** Min *has a p-well-formed ε-optimal strategy in $\mathcal{C}$.*

Constructing a classical strategy from a robust strategy preserving values. We now have all intermediate results to conclude the proof of Proposition 15. Let τ^p be an ε-optimal strategy for Max such that all its conforming plays are p-well-formed as defined in Lemma 16. The proof consists in the definition of an exact strategy, τ, from τ^p such that for every play ρ conforming to τ and starting from (ℓ,ν) with $\nu \notin \bigcup_{i\in\mathbb{N}^\star} F_i$, there exist a play ρ^r conforming to τ^p such that $\mathsf{wt}(\rho^r) \leq \mathsf{wt}(\rho)+\varepsilon$. Indeed, we obtain the following inequality:

$$\begin{aligned}\mathsf{rVal}^p_{\mathcal{C}}(\ell,\nu) &= \inf_{\sigma^r\in\mathsf{Strat}^p_{\mathcal{C},\mathsf{Min}}} \mathsf{wt}(\mathsf{rPlay}^p(\ell,\nu,\sigma^r,\tau^r))\\ &\leq \inf_{\sigma\in{\mathsf{Strat}^0}_{\mathcal{C},\mathsf{Min}}} \mathsf{wt}(\mathsf{Play}(\ell,\nu,\sigma,\tau))+\varepsilon = \mathsf{Val}_{\mathcal{C}}(\ell,\nu)+\varepsilon\end{aligned}$$

The first step to come up with τ is the definition of a "projection" mapping that transforms an exact play into a robust play. The main difficulty here is the need to adapt the decisions of Min when they are not robust (i.e. when Min plays "too close" to the upper-bound of a guard). We therefore use Lemma 17 to restrict the set of plays that we consider to plays conforming to a p-well-formed strategy. Formally, the projection $\mathsf{convert}^{\tau^p}_{wf\to r} : \mathsf{FPlays}_{wf} \to \mathsf{FPlays}^p$ is such that $\mathsf{convert}^{\tau^p}_{wf\to r}((\langle\ell,a\rangle,\nu)) = ((\langle\ell,a\rangle,\nu)$, and, for every $\rho=\rho_1 \xrightarrow{\delta,t} (\langle\ell,a\rangle,\nu)$, by letting ν_1 (resp. ν_1^c) be the last valuation of ρ_1 (resp. $\mathsf{convert}^{\tau^p}_{wf\to r}(\rho_1)$), we fix $\mathsf{convert}^{\tau^p}_{wf\to r}(\rho)$ to be equal to

- $\mathsf{convert}^{\tau^p}_{wf\to r}(\rho_1) \xrightarrow{\delta,\nu-\nu_1^c} s \xrightarrow{\tau^p(\mathsf{convert}^{\tau^p}_{wf\to r}(\rho_1)\xrightarrow{\delta,\nu-\nu_1^c} s)} s'$ if $\mathsf{last}(\rho_1)\in\mathsf{Conf}_{\mathsf{Min}}$
- $\mathsf{convert}^{\tau^p}_{wf\to r}(\rho_1) \xrightarrow{\tau^p(\mathsf{convert}^{\tau^p}_{wf\to r}(\rho_1))} s$ if $\mathsf{last}(\rho_1)\in\mathsf{Conf}_{\mathsf{Max}}$ and δ is chosen by τ^p
- $\mathsf{convert}^{\tau^p}_{wf\to r}(\rho_1) \xrightarrow{\delta,t} s$ otherwise

We note that, given ρ a p-well-formed play, $\mathsf{convert}^{\tau^p}_{wf\to r}(\rho)$ is indeed a robust play, since every decision for Max (resp. Min) on ρ is robust by Lemma 16 (resp. Lemma 17). Moreover, this projection enjoys the following interesting properties:

▶ **Lemma 18.** *Let ρ be a finite play that conforms to σ. Then,*
1. $\mathsf{convert}^{\tau^p}_{wf\to r}$ *preserves the last location of ρ. Formally, if $\langle\ell,a\rangle$ is the last location of ρ, then $\langle\ell,a\rangle$ is the last location of* $\mathsf{convert}_{wf\to r}(\rho)$*;*
2. $\mathsf{convert}^{\tau^p}_{wf\to r}$ *changes only slightly the last valuation of ρ. Formally, if ν is the last valuation of ρ, then the last valuation ν^c of* $\mathsf{convert}_{wf\to r}(\rho)$ *is such that $-N\cdot p \leq \nu^c-\nu \leq p$;*
3. $\mathsf{wt}_\Sigma(\mathsf{convert}_{wf\to r}(\rho)) \leq \mathsf{wt}_\Sigma(\rho) + W_{\mathsf{loc}}\cdot 2N\cdot p\cdot|\rho|$.

Using the projection mapping, we define the new strategy, τ, for Max to play in the exact semantics. Intuitively, τ mimicks as much as possible the choices of the robust strategy (by adapting delays to reach good valuations) against plays conforming to σ. Indeed, since σ

is ε-optimal for Min, it suffices for Max to play almost-optimally against it. Formally, for every (exact) play ρ conforming to σ, letting ν (resp. ν^c) be the last valuation of ρ (resp. $\mathsf{convert}^{\tau^p}_{wf\to r}(\rho)$), we define τ such that

$$\tau(\rho) = (\delta^{(j)}_{a\to b}, \nu^c + t^c - \nu) \qquad \text{if } (\delta^{(j)}_{a\to b}, t^c) = \tau^p(\mathsf{convert}^{\tau^p}_{wf\to r}(\rho))$$

On other play prefixes, τ may take any enable transition. So defined, τ is a legal strategy for Max, since the chosen delay satisfies the guard by reaching the same valuation of τ^p (that is a robust strategy). Moreover, by construction, we have the following property:

▶ **Lemma 19.** *If the play ρ conforms to τ and σ, then $\mathsf{convert}^{\tau^p}_{wf\to r}(\rho)$ conforms to τ^p.*

Given ρ a play conforming to τ and σ, we define $\rho' = \mathsf{convert}^{\tau^p}_{wf\to r}(\rho)$, and this play conforms to τ^p and satisfies $\mathsf{wt}(\rho') \leq \mathsf{wt}(\rho) + \varepsilon/2$. By Lemma 18, $\mathsf{wt}(\rho) \geq \mathsf{convert}^{\tau^p}_{wf\to r}(\rho) - \varepsilon$, Indeed, since $p \leq \eta$ and $|\rho| \leq \kappa$, we deduce that

$$\begin{aligned}
\mathsf{wt}_\Sigma(\rho) &\geq \mathsf{wt}_\Sigma(\mathsf{convert}^{\tau^p}_{wf\to r}(\rho)) - W_{\mathsf{loc}} \cdot 2N \cdot \eta \cdot \kappa \\
&> \mathsf{wt}_\Sigma(\mathsf{convert}^{\tau^p}_{wf\to r}(\rho)) - \frac{W_{\mathsf{loc}} \cdot 2N \cdot \kappa \cdot \varepsilon}{24N^2 \cdot W_{\mathsf{loc}} \cdot \kappa} \\
&\geq \mathsf{wt}_\Sigma(\mathsf{convert}^{\tau^p}_{wf\to r}(\rho)) - \varepsilon
\end{aligned}$$

From the observation that either ρ and $\mathsf{convert}^{\tau^p}_{wf\to r}(\rho)$ both reach the target or none of them do, and by by Lemma 18, we deduce that $\mathsf{wt}(\rho) \geq \mathsf{convert}^{\tau^p}_{wf\to r}(\rho) - \varepsilon$. That concludes the proof of Proposition 15.

6 Conclusion

This paper shows the decidability of the robust value problem for 1-clock weighted timed games under mild assumptions (no $-\infty$ value and a bounded clock $\nu \in [0, \mathsf{M}]$). We do so by reducing the robust value problem to the exact value problem in a WTG of exponential size (in the representation of the maximal constant M). Note that our proof technique more broadly establishes the computability of $\mathsf{rVal}^{0^+}(\ell_0, \nu)$, for any valuation $\nu \in [0, \mathsf{M}] \setminus \mathbb{N}^*$.

Our current decision procedure runs in doubly-exponential time. However, one can expect to lower this complexity: (1) instead of building all copies indexed from 0 to M, one could restrict to polynomially many copies indexed by the useful regions over 1-clock timed models [21], thus yielding a polynomial reduction to the exact value problem; (2) the exact value problem is suspected to be solvable in PSPACE (improving the known EXPTIME upper bound [23]) by an encoding into the existential theory of reals. If these two facts are confirmed, solving the value problem for 1-clock WTG with robustness would be in PSPACE, that is no harder (in terms of complexity classes) than the exact value problem for the same class: robustness induces no complexity increase.

As future work, beyond clarifying these complexity issues, we plan to explore whether our copy game construction extends to other classes of WTGs with decidable robust value problem, such as acyclic WTGs and divergent WTGs. If so, it would offer a tradeoff between the complexity (suspected higher with the copy game) and the freedom to choose the initial valuation (the copy game contruction currently only

References

1 Rajeev Alur, Mikhail Bernadsky, and P. Madhusudan. Optimal reachability for weighted timed games. In *Proceedings of the 31st International Colloquium on Automata, Languages and Programming (ICALP'04)*, volume 3142 of *Lecture Notes in Computer Science*, pages 122–133. Springer, 2004. doi:10.1007/978-3-540-27836-8_13.

2 Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994. doi:10.1016/0304-3975(94)90010-8.

3 Nathalie Bertrand, Maëlle Gautrin, and Julie Parreaux. Reaching as cheap as possible in 1-clock robust weighted timed games, 2026. arXiv:2606.28773.

4 Patricia Bouyer, Franck Cassez, Emmanuel Fleury, and Kim G. Larsen. Optimal strategies in priced timed game automata. In *Proceedings of the 24th International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'04)*, pages 148–160. Springer, 2004. doi:10.1007/978-3-540-30538-5_13.

5 Patricia Bouyer, Samy Jaziri, and Nicolas Markey. On the value problem in weighted timed games. In *Proceedings of the 26th International Conference on Concurrency Theory (CONCUR'15)*, volume 42 of *LIPIcs*, pages 311–324. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.CONCUR.2015.311.

6 Patricia Bouyer, Kim G Larsen, Nicolas Markey, and Jacob Illum Rasmussen. Almost optimal strategies in one clock priced timed games. In *Proceedings of the 26th International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'06)*, volume 4337 of *Lecture Notes in Computer Science*, pages 345–356. Springer, 2006. doi:10.1007/11944836_32.

7 Patricia Bouyer, Nicolas Markey, and Ocan Sankur. Robust weighted timed automata and games. In *Proceedings of the International Conference on Formal Modeling and Analysis of Timed Systems (Formats'13)*, volume 8053 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2013. doi:10.1007/978-3-642-40229-6_3.

8 Patricia Bouyer, Nicolas Markey, and Ocan Sankur. Robust reachability in timed automata and games: A game-based approach. *Theoretical Computer Science*, 563:43–74, 2015. doi:10.1016/J.TCS.2014.08.014.

9 Thomas Brihaye, Gilles Geeraerts, Axel Haddad, Engel Lefaucheux, and Benjamin Monmege. Simple priced timed games are not that simple. In *Proceedings of the 35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS'15)*, volume 45 of *LIPIcs*, pages 278–292. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.FSTTCS.2015.278.

10 Thomas Brihaye, Gilles Geeraerts, Axel Haddad, Engel Lefaucheux, and Benjamin Monmege. One-clock priced timed games with negative weights. *Logical Methods in Computer Science*, 18(3), 2022. doi:10.46298/lmcs-18(3:17)2022.

11 Thomas Brihaye, Gilles Geeraerts, Shankara Narayanan Krishna, Lakshmi Manasa, Benjamin Monmege, and Ashutosh Trivedi. Adding negative prices to priced timed games. In *Proceedings of the 25th International Conference on Concurrency Theory (CONCUR'14)*, volume 8704 of *Lecture Notes in Computer Science*, pages 560–575. Springer, 2014. doi:10.1007/978-3-662-44584-6_38.

12 Damien Busatto-Gaston, Benjamin Monmege, and Pierre-Alain Reynier. Optimal reachability in divergent weighted timed games. In *Proceedings of the 20th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'17)*, volume 10203 of *Lecture Notes in Computer Science*, pages 162–178, 2017. doi:10.1007/978-3-662-54458-7_10.

13 Damien Busatto-Gaston, Benjamin Monmege, and Pierre-Alain Reynier. Symbolic approximation of weighted timed games. In *Proceedings of the 38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'18)*, volume 122 of *LIPIcs*, pages 28:1–28:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.FSTTCS.2018.28.

14 Damien Busatto-Gaston, Benjamin Monmege, Pierre-Alain Reynier, and Ocan Sankur. Robust controller synthesis in timed Büchi automata: A symbolic approach. In *Proceedings of the 31st International Conference on Computer Aided Verification (CAV'19)*, volume 11561 of *Lecture Notes in Computer Science*, pages 572–590. Springer, 2019. doi:10.1007/978-3-030-25540-4_33.

15 John Fearnley, Rasmus Ibsen-Jensen, and Rahul Savani. One-clock priced timed games are PSPACE-hard. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'20)*, pages 397–409. ACM, 2020. doi:10.1145/3373718.3394772.

16 Shibashis Guha, Shankara Narayanan Krishna, Lakshmi Manasa, and Ashutosh Trivedi. Revisiting robustness in priced timed games. In *Proceedings of the 35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS'15)*, volume 45 of *LIPIcs*, pages 261–277. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.FSTTCS.2015.261.

17 Quentin Guilmant and Joël Ouaknine. Inaproximability in weighted timed games. In *Proceedings of the 35th International Conference on Concurrency Theory (CONCUR'24)*, volume 311 of *LIPIcs*, pages 27:1–27:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.CONCUR.2024.27.

18 Quentin Guilmant, Joël Ouaknine, and Isa Vialard. The value problem for weighted timed games with two clocks is undecidable. In *Proceedings of the 29th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'26)*, volume 16503 of *Lecture Notes in Computer Science*. Springer, 2026. doi:10.1007/978-3-032-22730-0_29.

19 Thomas Dueholm Hansen, Rasmus Ibsen-Jensen, and Peter Bro Miltersen. A faster algorithm for solving one-clock priced timed games. In *Proceedings of the 24th International Conference on Concurrency Theory (CONCUR'13)*, Lecture Notes in Computer Science, pages 531–545. Springer, 2013. doi:10.1007/978-3-642-40184-8_37.

20 Marcin Jurdzinski and Ashutosh Trivedi. Reachability-time games on timed automata. In *Proceedings of the 34th International Colloquium on Automata, Languages and Programming (ICALP'07)*, volume 4596 of *Lecture Notes in Computer Science*, pages 838–849. Springer, 2007. doi:10.1007/978-3-540-73420-8_72.

21 François Laroussinie, Nicolas Markey, and Philippe Schnoebelen. Model checking timed automata with one or two clocks. In *Proceedings of the 15th International Conference on Concurrency Theory (CONCUR'04)*, Lecture Notes in Computer Science, pages 387–401. Springer, 2004. doi:10.1007/978-3-540-28644-8_25.

22 Benjamin Monmege, Julie Parreaux, and Pierre-Alain Reynier. Synthesis of robust optimal real-time systems. In *Proceedings of the 49th International Symposium on Mathematical Foundations of Computer Science (MFCS'24)*, volume 306 of *LIPIcs*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.MFCS.2024.74.

23 Benjamin Monmege, Julie Parreaux, and Pierre-Alain Reynier. Decidability of one-clock weighted timed games with arbitrary weights. *Logical Methods in Computer Science*, 21(1), 2025. doi:10.46298/LMCS-21(1:8)2025.

24 Youssouf Oualhadj, Pierre-Alain Reynier, and Ocan Sankur. Probabilistic robust timed games. In *Proceedings of the 25th International Conference on Concurrency Theory (CONCUR'14)*, volume 8704 of *Lecture Notes in Computer Science*, pages 203–217. Springer, 2014. doi:10.1007/978-3-662-44584-6_15.

25 Anuj Puri. Dynamical properties of timed automata. *Discrete Event Dynamical Systems*, 10(1-2):87–113, 2000. doi:10.1023/A:1008387132377.

26 Ocan Sankur, Patricia Bouyer, and Nicolas Markey. Shrinking timed automata. *Information and Computation*, 234:107–132, 2014. doi:10.1016/J.IC.2014.01.002.

27 Ocan Sankur, Patricia Bouyer, Nicolas Markey, and Pierre-Alain Reynier. Robust controller synthesis in timed automata. In *Proceedings of the 24th International Conference on Concurrency Theory (CONCUR'13)*, volume 8052 of *Lecture Notes in Computer Science*, pages 546–560. Springer, 2013. doi:10.1007/978-3-642-40184-8_38.

28 Martin De Wulf, Laurent Doyen, Nicolas Markey, and Jean-François Raskin. Robustness and implementability of timed automata. In *Proceedings of the Joint International Conferences on Formal Modelling and Analysis of Timed Systems (Formats'04) and Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'04)*, volume 3253 of *Lecture Notes in Computer Science*, pages 118–133. Springer, 2004. doi:10.1007/978-3-540-30206-3_10.