Synthesis of Robust Optimal Real-Time Systems

Benjamin Monmege
Aix Marseille Univ, CNRS, LIS, Marseille, France

Julie Parreaux
University of Warsaw, Poland

Pierre-Alain Reynier
Aix Marseille Univ, CNRS, LIS, Marseille, France

Abstract

Weighted Timed Games (WTGs for short) are widely used to describe real-time controller synthesis problems, but they rely on an unrealistic perfect measure of time elapse. In order to produce strategies tolerant to timing imprecisions, we consider a notion of robustness, expressed as a parametric semantics, first introduced for timed automata. WTGs are two-player zero-sum games played in a weighted timed automaton in which one of the players, that we call Min, wants to reach a target location while minimising the cumulated weight. The opponent player, in addition to controlling some of the locations, can perturb delays chosen by Min. The robust value problem asks, given some threshold, whether there exists a positive perturbation and a strategy for Min ensuring to reach the target, with an accumulated weight below the threshold, whatever the opponent does.

We provide in this article the first decidability result for this robust value problem. More precisely, we show that we can compute the robust value function, in a parametric way, for the class of divergent WTGs (this class has been introduced previously to obtain decidability of the (classical) value problem in WTGs without bounding the number of clocks). To this end, we show that the robust value is the fixpoint of some operators, as is classically done for value iteration algorithms. We then combine in a very careful way two representations: piecewise affine functions introduced in [1] to analyse WTGs, and shrunk Difference Bound Matrices (shrunk DBMs for short) considered in [29] to analyse robustness in timed automata. The crux of our result consists in showing that using this representation, the operator of value iteration can be computed for infinitesimally small perturbations. Last, we also study qualitative decision problems and close an open problem on robust reachability, showing it is EXPTIME-complete for general WTGs.

2012 ACM Subject Classification Software and its engineering → Formal software verification; Theory of computation → Algorithmic game theory

Keywords and phrases Weighted timed games, Algorithmic game theory, Robustness

Digital Object Identifier 10.4230/LIPIcs.MFCS.2024.74

Related Version *Full Version*: https://arxiv.org/abs/2403.06921 [25]

Funding This work has been partly funded by the QuaSy project (ANR-23-CE48-0008) and the NCN grant 2019/35/B/ST6/02322.

1 Introduction

The design and synthesis of real-time systems have long been paramount challenges, given the critical need for dependable and efficient systems in a variety of applications. In particular, the pursuit of robustness and reliability in these systems has led researchers to explore innovative methods and formalisms to address the complexities inherent in real-time environments. In this work, we focus on game-based models, and more precisely on the game extension of timed automata [2], a.k.a. timed games, which provide an elegant framework for capturing the interplay between system components, environment dynamics, and strategic decision-making. More precisely, in this model, locations of a timed automaton are split amongst the two players, which play in turn in the infinite-state space of the automaton.

Regarding robustness, prior studies have primarily focused on areas such as fault tolerance, adaptive control, and formal methods. In this work, we follow a series of works based on game theory. The objective is to fill the gap between mathematical models such as timed automata, often used for model-checking purposes, and implementation constraints, in which clocks only have finite precision, and actions are not instantaneous. To that end, a parametric semantics has been considered in [27], which consists in allowing the delays to be perturbed by some limited amount. The uncertainty of the system, i.e. the perturbation of the delays, is modelled by an adversarial environment. Two kinds of problems can then be considered: first, the analysis may be done for a fixed perturbation bound (we call it a *fixed-perturbation robustness problem*); second, in order to abstract the precise settings of the implementation, and as the exact value of the perturbation bound may be unknown, one can try to determine whether there exists a perturbation bound under which the system is reliable (we call it an *existential robustness problem*). By monotonicity of the semantics w.r.t. the perturbation bound, if one manages to prove the reliability against some perturbation bound, then it still holds for smaller perturbations. Initially introduced for model-checking purposes [20, 8], this approach has been lifted to automatic synthesis, yielding the so-called *conservative semantics*, studied for instance in [30, 26, 19]. In these works, a player named Controller aims at satisfying a liveness objective while its opponent may perturb the delays.

In the present work, we aim to go beyond qualitative objectives, and tackle quantitative aspects. In real-time systems and critical applications, quantitative aspects such as resource utilization and performance metrics hold important significance. This has led to the model of *weighted timed games* (WTG for short), which has been widely studied during the last two decades. When considering a reachability objective, Controller (a.k.a. player Min) aims at reaching a set of target locations while minimizing the accumulated weight. One is then interested in the *value problem*, which consists in deciding, given some threshold, whether a strategy for Min exists to reach some target location while keeping the accumulated cost below this threshold. While this problem is undecidable in general [11, 5, 14], several subclasses have been identified that allow one to regain decidability. Amongst recent works, we can cite the class of divergent WTGs [18] which generalize to arbitrary costs the class of strictly non-Zeno costs introduced in [6], or the class of one-clock WTGs [24].

The core objective of this research is to explore the synthesis of real-time systems that not only meet timing constraints but also optimize performance with respect to specified weight objectives and are robust against timing imprecisions. To that end, we aim to study the setting of timed games extended with both robustness issues and quantitative aspects. We focus on the conservative semantics and on reachability objectives. In this setting, under a fixed perturbation, the player Min aims at reaching a set of target locations while minimizing the accumulated weight, and resisting delay perturbations. This leads to a notion of *robust value under a fixed perturbation*: this is simply the best value Min can achieve. The associated fixed perturbation robust value problem aims at comparing this value with a given threshold. When turning to the existential robustness decision problem, one considers the notion of *robust value*, defined as the limit of robust values for arbitrarily small perturbation values. We prove that this limit exists, and study the associated decision problem, which we simply call *robust value problem*, and which can be defined as follows: given a threshold, determine whether there exists a positive perturbation, and a winning strategy for Min ensuring that the accumulated weight until the target is below the threshold.

This problem is highly challenging as it combines difficulties coming from the introduction of weights, with those due to the analysis of an existential problem for the parametric semantics of robustness. Unsurprisingly, it has been shown to be undecidable [28]. To highlight the challenges we face, already in the qualitative setting (w/o weights), the setting

of two-player has not been addressed yet for the conservative semantics, hence the existential robust reachability problem was left open in [28]. Indeed, in [19, 30], only the one-player case is handled (a partial extension is considered in [26]) and the existential robustness reachability problem for the two-player setting has only been solved for the excessive semantics (an alternative to the conservative semantics) in [10]. Regarding the quantitative setting, very few works have addressed robustness issues. The fixed-perturbation robust value problem is shown to be decidable for one-clock weighted timed games in [21], with non-negative weights only, and for the excessive semantics. In [9], the authors consider the one-player case and prove that the robust value problem is PSPACE-complete.

Our contributions are as follows: first, regarding the qualitative setting, we close the case of existential robust reachability in two-player timed games for the conservative semantics, and show that this problem is EXPTIME-complete. To do so, we introduce a construction which allows us to reduce the problem to the excessive semantics solved in [10]. As a corollary, we deduce an upper bound on the length of paths to the target.

Then, we turn to the quantitative setting and show that for the class of divergent WTGs (one of the largest classes of WTGs for which the decidability of the value problem is known), the robust value problem is decidable. We proceed as follows:

1. We characterize the robust value for a fixed perturbation as the fixpoint of some operator.
2. We show that for acyclic WTGs, this fixpoint can be obtained as a finite iteration of this operator, which we decompose using four simpler operators.
3. We introduce a symbolic parametric approach for the computation of this operator, for arbitrarily small values of the perturbation. This requires carefully combining the representation of value functions using piecewise affine functions introduced in [1] with the notion of shrunk DBMs, used in [29] to analyse robustness issues in timed automata. This yields the decidability of the robust value problem for the class of acyclic WTGs.
4. By combining this with the upper bound deduced from the qualitative analysis, we show the decidability of the robust value problem for the whole class of divergent WTGs.

In Section 2, we introduce WTGs, under the prism of robustness. We describe in Section 3 the robustness problems we consider, present our contributions for qualitative ones, and state that we can solve the quantitative one for acyclic WTGs. Sections 4 and 5 detail how to prove this result, following steps 1.-3. described above. Last, Section 6 extends this positive result to the class of divergent WTGs. Omitted proofs can be found in a long version of this article [25].

2 Robustness in weighted timed games

We let $\mathcal{X}$ be a finite set of variables called clocks. A valuation is a mapping $\nu\colon \mathcal{X} \to \mathbb{R}_{\geq 0}$. For a valuation ν, a delay $t \in \mathbb{R}_{\geq 0}$ and a subset $Y \subseteq \mathcal{X}$ of clocks, we define the valuation $\nu + t$ as $(\nu + t)(x) = \nu(x) + t$, for all $x \in \mathcal{X}$, and the valuation $\nu[Y := 0]$ as $(\nu[Y := 0])(x) = 0$ if $x \in Y$, and $(\nu[Y := 0])(x) = \nu(x)$ otherwise. A (non-diagonal) guard on clocks of $\mathcal{X}$ is a conjunction of atomic constraints of the form $x \bowtie c$, where $\bowtie \in \{\leq, <, =, >, \geq\}$ and $c \in \mathbb{N}$. A valuation $\nu\colon \mathcal{X} \to \mathbb{R}_{\geq 0}$ satisfies an atomic constraint $x \bowtie c$ if $\nu(x) \bowtie c$. The satisfaction relation is extended to all guards g naturally, and denoted by $\nu \models g$. We let $\mathsf{Guards}(\mathcal{X})$ denote the set of guards over $\mathcal{X}$.

▶ **Definition 1.** *A* weighted timed game *(WTG) is a tuple* $\mathcal{G} = \langle L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_T, \mathcal{X}, \Delta, \mathsf{wt}\rangle$ *where* L_{Min}*,* L_{Max}*,* L_T *are finite disjoint subsets of* Min *locations,* Max *locations, and target locations, respectively (we let* $L = L_{\mathsf{Min}} \uplus L_{\mathsf{Max}} \uplus L_T$*),* $\mathcal{X}$ *is a finite set of clocks,* $\Delta \subseteq L \times \mathsf{Guards}(\mathcal{X}) \times 2^{\mathcal{X}} \times L$ *is a finite set of transitions, and* $\mathsf{wt}\colon \Delta \uplus L \to \mathbb{Z}$ *is the weight function.*

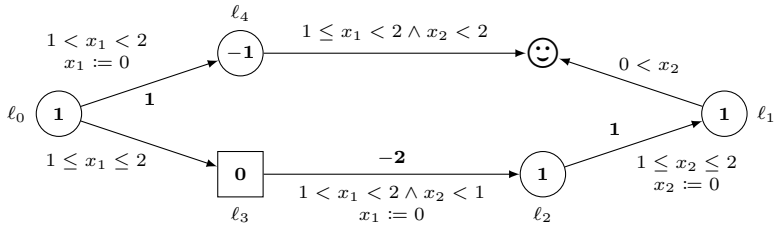

Figure 1 An acyclic WTG with two clocks.

The usual semantics, called *exact semantics*, of a WTG $\mathcal{G}$ is defined in terms of a game played on an infinite transition system whose vertices are configurations of the WTG denoted by $\mathsf{Conf} = \mathsf{Conf}_{\mathsf{Min}} \uplus \mathsf{Conf}_{\mathsf{Max}} \uplus \mathsf{Conf}_T$. A configuration is a pair (ℓ, ν) with a location and a valuation of the clocks. A configuration is final (resp. belongs to Min, or Max), and belongs to Conf_T (resp. to $\mathsf{Conf}_{\mathsf{Min}}$, or $\mathsf{Conf}_{\mathsf{Max}}$) if its location is a target location of L_T (resp. of L_{Min}, or L_{Max}). The alphabet of the transition system is given by $\Delta \times \mathbb{R}_{\ge 0}$: a pair (δ, t) encodes the delay t that a player wants to spend in the current location, before firing transition δ. An example of WTG is depicted in Figure 1.

In this article, we consider an alternative semantics to model the robustness, traditionnally called the *conservative semantics*. It is defined in a WTG $\mathcal{G}$ according to a fixed parameter $p > 0$. This semantics allows Max to slightly perturb the delays chosen by Min with an amplitude bounded by p. From the modelling perspective, the perturbations model the small errors of physical systems on the real value of clocks. Conservative means that the delays proposed by Min must remain feasible after applying all possible perturbations. In particular, the conservative semantics does not add new edges with respect to the exact one.

▶ **Definition 2.** *Let $\mathcal{G} = \langle L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_T, \mathcal{X}, \Delta, \mathsf{wt}\rangle$ be a WTG. For $p \ge 0$, we let $[\![\mathcal{G}]\!]^p = \langle S, E, \mathsf{wt}\rangle$ with $S = S_{\mathsf{Min}} \uplus S_{\mathsf{Max}} \uplus S_T$ the set of* states *with $S_{\mathsf{Min}} = \mathsf{Conf}_{\mathsf{Min}}$, $S_T = \mathsf{Conf}_T$ and $S_{\mathsf{Max}} = \mathsf{Conf}_{\mathsf{Max}} \cup (\mathsf{Conf}_{\mathsf{Min}} \times \mathbb{R}_{\ge 0} \times \Delta)$; $E = E_{\mathsf{Min}} \uplus E_{\mathsf{Max}} \uplus E_{\mathit{rob}}$ the set of* edges *with*

$$E_{\mathsf{Max}} = \left\{ \left((\ell,\nu) \xrightarrow{\delta,t} (\ell',\nu')\right) \mid \ell \in L_{\mathsf{Max}}, \nu + t \models g \text{ and } \nu' = (\nu + t)[Y := 0] \right\}$$

$$E_{\mathsf{Min}} = \left\{ \left((\ell,\nu) \xrightarrow{\delta,t} (\ell,\nu,\delta,t)\right) \mid \ell \in L_{\mathsf{Min}}, \nu + t \models g \text{ and } \nu + t + 2p \models g \right\}$$

$$E_{\mathit{rob}} = \left\{ \left((\ell,\nu,\delta,t) \xrightarrow{\delta,\varepsilon} (\ell',\nu')\right) \mid \varepsilon \in [0,2p] \text{ and } \nu' = (\nu + t + \varepsilon)[Y := 0] \right\}$$

where $\delta = (\ell, g, Y, \ell') \in \Delta$; and $\mathsf{wt}\colon S \cup E \to \mathbb{Z}$ 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 $e \in E$, $\mathsf{wt}(e) = \mathsf{wt}(\delta)$ if $e = (s \xrightarrow{\delta,t} s')$ with $s \in \mathsf{Conf}$, or $\mathsf{wt}(e) = 0$ otherwise.

When $p = 0$, the infinite transition system $[\![\mathcal{G}]\!]^0$ describes the exact semantics of the game, the usual semantics where each step of the player Min is cut into the true step, followed by a useless edge $(\ell,\nu,\delta,t) \xrightarrow{\delta,0} (\ell',\nu')$ where Max has no choice. When $p > 0$, the infinite transition system $[\![\mathcal{G}]\!]^p$ describes the conservative semantics of the game: states $(\ell,\nu,\delta,t) \in \mathsf{Conf}_{\mathsf{Min}} \times \mathbb{R}_{\ge 0} \times \Delta$ where Max must choose the perturbation in the interval $[0, 2p]$ are called *perturbed states*.

Let s be a state of $[\![\mathcal{G}]\!]^p$, we denote by $E(s)$ the set of possible outgoing edges of $[\![\mathcal{G}]\!]^p$ from s. We extend this notation to locations to denote the set of outgoing transitions in $\mathcal{G}$. A state (resp. location) s is a *deadlock* when $E(s) = \emptyset$. We note that the conservative semantics may introduce deadlock in configurations of Min (even if an outgoing edge exists in the exact semantics). Thus, unlike in the literature [1, 18], we allow state *and* location deadlocks.

A *finite play* of $\mathcal{G}$ w.r.t. the conservative semantics with parameter p is a sequence of edges in the transition system $[\![\mathcal{G}]\!]^p$ starting in a configuration of $\mathcal{G}$. We denote by $|\rho|$ the number of edges of ρ, and by $\mathsf{last}(\rho)$ its last state. The concatenation of two finite plays ρ_1 and ρ_2, such that ρ_1 ends in the same state as ρ_2 starts, is denoted by $\rho_1\rho_2$. Moreover, for modelling reasons, we only consider finite plays (starting and) ending in a configuration of $\mathcal{G}$. Since a finite play is always defined regarding a parameter p for the conservative semantics, we denote by FPlays^p this set of finite plays. Moreover, we denote by $\mathsf{FPlays}^p_{\mathsf{Min}}$ (resp. $\mathsf{FPlays}^p_{\mathsf{Max}}$) the subset of these finite plays ending in a state of Min (resp. Max). A *maximal play* is then a maximal sequence of consecutive edges: it is either a finite play reaching a deadlock (not necessary in L_T), or an infinite sequence such that all its prefixes are finite plays.

The objective of Min is to reach a target configuration, while minimising the cumulated weight up to the target. Hence, we associate to every finite play $\rho = s_0 \xrightarrow{\delta_0,t_0} s_1 \xrightarrow{\delta_1,t_1} \cdots s_k$ (some edges are in E_{rob}, others are not) its cumulated weight, taking into account both discrete and continuous costs: $\mathsf{wt}_\Sigma(\rho) = \sum_{i=0}^{k-1}[t_i \times \mathsf{wt}(s_i) + \mathsf{wt}(\delta_i)]$. Then, the weight of a maximal play ρ, denoted by $\mathsf{wt}(\rho)$, is defined by $+\infty$ if ρ does not reach L_T (because it is infinite or reaches another deadlock), and $\mathsf{wt}_\Sigma(\rho)$ if it ends in (ℓ,ν) with $\ell \in L_T$.

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. Since plays could reach a deadlock state of Min, 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_2,ν) can be defined only when $\nu(x_2) \leq 2-2p$ since, otherwise, there are no outgoing edges in $[\![\mathcal{G}]\!]^p$ from this state. Symmetrically, we ask for Max to always propose a move if we are not in a deadlock state. More formally, a strategy for Min, denoted χ, is a (possibly partial) mapping $\chi\colon \mathsf{FPlays}^p_{\mathsf{Min}} \to E$ such that $\chi(\rho) \in E(\mathsf{last}(\rho))$. A strategy for Max, denoted ζ, is a (possibly partial) mapping $\zeta\colon \mathsf{FPlays}^p_{\mathsf{Max}} \to E$ such that for all ρ, if $E(\mathsf{last}(\rho)) \neq \emptyset$, then $\chi(\rho)$ is defined, and in this case belongs to $E(\mathsf{last}(\rho))$. The set of strategies of Min (resp. Max) with the perturbation p is denoted by $\mathsf{Strat}^p_{\mathsf{Min}}$ (resp. $\mathsf{Strat}^p_{\mathsf{Max}}$).

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 χ of Min (resp. Max) if for all k such that s_k belongs to Min (resp. Max), we have that $(\delta_k,t_k) = \chi(s_0 \xrightarrow{\delta_0,t_0} \cdots s_k)$. For all strategies χ and ζ of players Min and Max, respectively, and for all configurations (ℓ_0,ν_0), we let $\mathsf{Play}((\ell_0,\nu_0),\chi,\zeta)$ be the outcome of χ and ζ, defined as the unique maximal play conforming to χ and ζ and starting in (ℓ_0,ν_0).

The semantics $[\![\mathcal{G}]\!]^p$ is monotonic with respect to the perturbation p in the sense that Min has more strategies when p decreases, while Max can obtain, against a fixed strategy of Min, a smaller weight when p decreases. Formally, we have:

▶ **Lemma 3.** *Let $\mathcal{G}$ be a WTG, and $p > p' \geq 0$ be two perturbations. Then*

1. $\mathsf{Strat}^p_{\mathsf{Min}} \subseteq \mathsf{Strat}^{p'}_{\mathsf{Min}}$;
2. *for all* $\chi \in \mathsf{Strat}^p_{\mathsf{Min}}$, $\sup_{\zeta\in\mathsf{Strat}^p_{\mathsf{Max}}} \mathsf{wt}(\mathsf{Play}((\ell,\nu),\chi,\zeta)) \geq \sup_{\zeta\in\mathsf{Strat}^{p'}_{\mathsf{Max}}} \mathsf{wt}(\mathsf{Play}((\ell,\nu),\chi,\zeta))$.

3 Deciding the robustness in weighted timed games

We aim to study what Min can guarantee, qualitatively and then quantitatively, in the conservative semantics of weighted timed games whatever Max does.

Qualitative robustness problems. Formally, given a WTG $\mathcal{G}$ and a perturbation p, we say a strategy χ of Min is *winning* in $[\![\mathcal{G}]\!]^p$ from configuration (ℓ,ν) if for all strategies ζ of Max, $\mathsf{Play}((\ell,\nu),\chi,\zeta)$ is a finite play ending in a location of L_T.

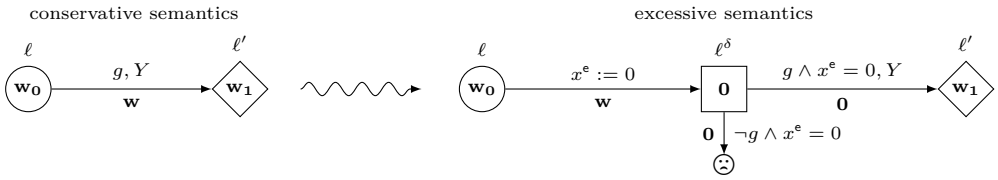

Figure 2 Gadget used to encode the conservative semantics into the excessive one. Each transition $\delta = (\ell, g, Y, \ell')$ with $\ell \in L_{\mathsf{Min}}$ is replaced by the gadget. Symbols w, w_0, w_1 denote weights from $\mathcal{G}$. The new location ℓ^δ of Max uses a fresh clock x^{e} to test the guard after the perturbation (as in the conservative semantics). The new location ☹ is a deadlock (thus winning for Max).

There are two possible questions, whether the perturbation p is fixed, or if we should consider it to be infinitesimally small:

- *fixed-perturbation robust reachability problem*: given a WTG $\mathcal{G}$, a configuration (ℓ, ν) and a perturbation $p > 0$, decide whether Min has a winning strategy χ from (ℓ, ν) in $[\![\mathcal{G}]\!]^p$;
- *existential robust reachability problem*: given a WTG $\mathcal{G}$ and a configuration (ℓ, ν), decide whether there exists $p > 0$ such that Min has a winning strategy χ from (ℓ, ν) in $[\![\mathcal{G}]\!]^p$. Notice that by Lemma 3, if Min has a winning strategy χ from (ℓ, ν) in $[\![\mathcal{G}]\!]^p$, then he has one in $[\![\mathcal{G}]\!]^{p'}$ for all $p' \leq p$.

When the perturbation p is fixed, we can encode in a WTG the conservative semantics described in $[\![\mathcal{G}]\!]^p$, by adding new locations for Max to choose a perturbation, and by modifying the guards that will now use the perturbation p. Solving the fixed-perturbation robust reachability problem then amounts to solving a reachability problem in the modified WTG[1] which can be performed in EXPTIME [3] (here weights are useless). Since the reachability problem in timed games is already EXPTIME-complete [22], we obtain:

▶ **Proposition 4.** *The fixed-perturbation robust reachability problem is* EXPTIME*-complete.*

We now turn our attention to the existential robust reachability problem. This problem was left open for the conservative semantics (see [28], Table 1.2 page 17), while it has been solved in [10] for an alternative semantics of robustness, known as the *excessive semantics*. Intuitively, while the conservative semantics requires that the delay, after perturbation, satisfies the guard, the excessive semantics only requires that the delay, *without perturbation*, satisfies the guard. We present a reduction from the conservative semantics to the excessive one, allowing us to solve the existential robust reachability problem for the conservative semantics. Intuitively, the construction (depicted on Figure 2) adds a new location (for Max) for each transitions of Min to test the delay chosen by Min after the perturbation:

▶ **Proposition 5.** *The existential robust reachability problem is* EXPTIME*-complete.*

Quantitative fixed-perturbation robustness problem. We are also interested in the minimal weight that Min can guarantee while reaching the target whatever Max does: to do that we define *robust values*. First, we define the *fixed-perturbation robust value*: for all configurations (ℓ, ν) of $\mathcal{G}$ (and not for all states of the semantics), we let $\overline{\mathsf{rVal}}^p(\ell, \nu) = \inf_{\chi \in \mathsf{Strat}^p_{\mathsf{Min}}} \sup_{\zeta \in \mathsf{Strat}^p_{\mathsf{Max}}} \mathsf{wt}(\mathsf{Play}((\ell, \nu), \chi, \zeta))$.

[1] By transforming the WTG, its guards use rational instead of natural numbers (due to p). To fit the classical definition of WTG, we can apply a scaling factor (i.e. $1/p$) to all constants appearing in this WTG. We note that this operation preserves the set of winning strategies for the reachability objective (here weights are irrelevant) by applying the scaling operations on strategies too.

Since a fixed-perturbation conservative semantics defines a *quantitative reachability game*[2], where configurations of Max also contain the robust states, we obtain that the fixed-perturbation robust value is determined, by applying [13, Theorem 2.2], i.e. that $\overline{\mathsf{rVal}}^p(\ell,\nu) = \sup_{\zeta\in\mathsf{Strat}^p_{\mathsf{Max}}} \inf_{\chi\in\mathsf{Strat}^p_{\mathsf{Min}}} \mathsf{wt}(\mathsf{Play}((\ell,\nu),\chi,\zeta))$. We therefore denote rVal^p this value.

▶ Remark 6. In [9, 10, 21], the set of possible perturbations for Max is $[-p,p]$. For technical reasons, we use a (equivalent) perturbation with a shift of the delay proposed by Min by p.

When $p=0$, rVal^0 defines the *(exact) value* that is used to study the value problem in WTGs. By Lemma 3, we can deduce that the fixed-perturbation robust value is monotonic with respect to the perturbation p and is always an upper-bound for the (exact) value.

▶ **Lemma 7.** *Let $\mathcal{G}$ be a WTG, and $p > p' \geq 0$ be two perturbations. Then, for all configurations (ℓ,ν), $\mathsf{rVal}^p(\ell,\nu) \geq \mathsf{rVal}^{p'}(\ell,\nu)$.*

As in the qualitative case, when the perturbation p is fixed, we can encode in polynomial time in a WTG the conservative semantics described in $[\![\mathcal{G}]\!]^p$. Unfortunately, the value of WTGs is not always computable since the associated decision problems (in particular the *value problem* that requires to decide if the value of a given configuration is below a given threshold) are undecidable [11, 4, 14]. However, in subclasses of WTGs where the value function can be computed, like acyclic WTGs (where every path in the graph of the WTGs is acyclic, decidable in 2-EXPTIME [15]) or divergent WTGs (that we recall in Section 6, in 3-EXPTIME [6, 16]), the fixed-perturbation robust value is also computable (if the modified game falls in the subclass). In particular, we obtain:

▶ **Proposition 8.** *We can compute the fixed-perturbation robust value of a WTG that is acyclic (in 2-EXPTIME) or divergent (in 3-EXPTIME), for all possible initial configurations.*

On top of computing the robust values, the previous works also allow one to synthesize almost-optimal (*i.e.* arbitrarily close from the value) strategies for both players.

Quantitative robustness problem. Now, we consider the existential version of this problem by considering an infinitesimal perturbation. We thus want to know what Min can guarantee as a value if Max plays infinitesimally small perturbations. To define properly the problem, we introduce a new value: given a WTG $\mathcal{G}$, the *robust value* is defined, for all configurations (ℓ,ν) of $\mathcal{G}$, by $\overline{\mathsf{rVal}}(\ell,\nu) = \lim_{p\to 0, p>0} \mathsf{rVal}^p(\ell,\nu)$. This value is defined as a limit of functions (the fixed-perturbation robust values), which can be proved to always exist as the limit of a non-increasing sequence of functions (see Lemma 7). The decision problem associated to this robust value is: given a WTG $\mathcal{G}$, an initial configuration (ℓ,ν), and a threshold $\lambda \in \mathbb{Q}\cup\{-\infty,+\infty\}$, decide if $\overline{\mathsf{rVal}}(\ell,\nu) \leq \lambda$. We call it the *robust value problem*.

Unsurprisingly, this problem is undecidable [9, Theorem 4]. We will thus consider some restrictions over WTGs. In particular, we consider classes of WTGs where the (non robust) value problem is known as decidable for the exact semantics: acyclic WTGs [1], and divergent WTGs [6, 16]. Our first main contribution concerns the acyclic case:

▶ **Theorem 9.** *The robust value problem is decidable over the subclass of acyclic WTGs.*

The next two sections sketch the proof of this theorem via an adaptation of the value iteration algorithm [1] used to compute the value function in (non-robust) acyclic WTGs: it consists in computing iteratively the best thing that both players can guarantee in a bounded

[2] A quantitative reachability game introduced in [12] is an abstract model to formally define the semantics of quantitative (infinite) games.

number of steps, that increases step by step. It is best described by a mapping $\mathcal{F}$ that explains how a value function gets modified by allowing one more step for the players to play. The adaptation we propose consists in taking the robustness into account by using shrunk DBM techniques introduced in [29]: instead of inequalities on the difference of two clock values of the form $x - y \leq c$ involving rational constants c, the constants c will now be of the form $a - bp$, with a a rational and b a positive integer, p being an infinitesimal perturbation. This will allow us to compute a description of the fixed-perturbation values for all initial configurations, for all perturbations p smaller than an upperbound that we will compute. The robust value will then be obtained by taking the limit of this parametric representation of the fixed-perturbation values when p tends to 0. Our algorithm will also compute an upperbound on the biggest allowed perturbation p. As in previous works, once the value is computed, we can also synthesize almost-optimal strategies.

Section 4 will describe the mapping $\mathcal{F}_p$, with a known perturbation p: the iteration of this operator will be shown to converge towards the fixed-perturbation value rVal^p. The robust value functions will be shown to always be piecewise affine functions with polytope pieces. Section 5 describes the parametric representation of these functions, where the perturbation is no longer fixed but is a formal parameter $\mathfrak{p}$. We then explain how the mapping $\mathcal{F}_p$ can be computed for all small enough values of p at once, allowing us to conclude.

4 Operator $\mathcal{F}_p$ to compute the fixed-perturbation value

The first step of the proof is the definition of the new operator adapted from the operator $\mathcal{F}$ of [1]. We thus fix a perturbation $p > 0$, and we define an operator $\mathcal{F}_p$ taking as input a mapping $X \colon L \times \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$, computing a mapping $\mathcal{F}_p(X) \colon L \times \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$ defined for all configurations (ℓ, ν) by $\mathcal{F}_p(X)(\ell, \nu)$ equal to

$$\begin{cases} 0 & \text{if } \ell \in L_T \\ +\infty & \text{if } \ell \in L_{\mathsf{Max}} \text{ and } E(\ell,\nu) = \emptyset \quad \text{(if } \mathsf{Max} \text{ reaches a deadlock, he wins)} \\ \sup_{(\ell,\nu) \xrightarrow{\delta,t} (\ell',\nu') \in [\![\mathcal{G}]\!]^p} \left[t\,\mathsf{wt}(\ell) + \mathsf{wt}(\delta) + X(\ell',\nu') \right] & \text{if } \ell \in L_{\mathsf{Max}} \text{ and } E(\ell,\nu) \neq \emptyset \\ \inf_{(\ell,\nu) \xrightarrow{\delta,t} (\ell,\nu,\delta,t) \in [\![\mathcal{G}]\!]^p} \sup_{(\ell,\nu,\delta,t) \xrightarrow{\delta,\varepsilon} (\ell,\nu') \in [\![\mathcal{G}]\!]^p} \left[(t+\varepsilon)\,\mathsf{wt}(\ell) + \mathsf{wt}(\delta) + X(\ell',\nu') \right] & \text{if } \ell \in L_{\mathsf{Min}} \end{cases}$$

In the following, we let $\mathbb{V}^0$ be the mapping $L \times \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$ defined by $\mathbb{V}^0(\ell,\nu) = 0$ if $\ell \in L_T$ and $\mathbb{V}^0(\ell,\nu) = +\infty$ otherwise. By adapting the proof of the non-robust setting, we show:

▶ **Lemma 10.** *Let $\mathcal{G}$ be an acyclic WTG, $p > 0$, D is the depth of $\mathcal{G}$, i.e. the length of a longest path in $\mathcal{G}$. Then, rVal^p is a fixpoint of $\mathcal{F}_p$, and $\mathsf{rVal}^p = \mathcal{F}_p^D(\mathbb{V}^0)$.*

We can thus compute the fixed-perturbation robust value of an acyclic WTG by repeatedly computing $\mathcal{F}_p$. We will see in the next section that this computation can be made for all small enough p by using a parametric representation of the mappings. It will be easier to split the computation of $\mathcal{F}_p$ in several steps (as done in the non robust case [1, 18]). Each of the four operators takes as input a mapping $V \colon \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$ (where the location ℓ has been fixed, with respect to mappings $L \times \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$), and computes a mapping of the same type.

- The operator $\mathsf{Unreset}_Y$, with $Y \subseteq \mathcal{X}$ a subset of clocks, is such that for all $\nu \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$, $\mathsf{Unreset}_Y(V)(\nu) = V(\nu[Y := 0])$.
- The operator Guard_δ, with $\delta = (\ell, g, Y, \ell')$ a transition of Δ, is such that for all $\nu \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$, if $\nu \models g$, then $\mathsf{Guard}_\delta(V)(\nu) = V(\nu)$; otherwise, $\mathsf{Guard}_\delta(V)(\nu)$ is equal to $-\infty$ if $\ell \in L_{\mathsf{Max}}$, and $+\infty$ if $\ell \in L_{\mathsf{Min}}$.

- The operator Pre_ℓ, with $\ell \in L$, is such that for all $\nu \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$, $\mathsf{Pre}_\ell(V)(\nu)$ is equal to $\sup_{t\geq 0}[t\,\mathsf{wt}(\ell) + V(\nu+t)]$ if $\ell \in L_{\mathsf{Max}}$, and $\inf_{t\geq 0}[t\,\mathsf{wt}(\ell) + V(\nu+t)]$ if $\ell \in L_{\mathsf{Min}}$.
- The operator $\mathsf{Perturb}^p_\ell$, with perturbation $p > 0$ and $\ell \in L_{\mathsf{Min}}$, is such that for all $\nu \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$, $\mathsf{Perturb}^p_\ell(V)(\nu) = \sup_{\varepsilon\in[0,2p]}[\varepsilon\,\mathsf{wt}(\ell) + V(\nu+\varepsilon)]$.

Though the situation is not symmetrical for Min and Max in $\mathcal{F}_p$, in particular for the choice of delay t, the definition of Pre_ℓ does not differentiate the two players with respect to their choice of delay. However, the correctness of the decomposition comes from the combination of this operator with Guard_δ (that clearly penalises Min if he chooses a delay such that the translated valuation does not satisfy the guard) and $\mathsf{Perturb}^p_\ell$ that allows Max to select a legal perturbation not satisfying the guard, leading to a value $+\infty$ afterwards.

For a mapping $\mathbb{V}\colon L \times \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$ and a location ℓ, we can extract the submapping for the location ℓ, that we denote by $\mathbb{V}_\ell\colon \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$, defined for all $\nu \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$ by $\mathbb{V}_\ell(\nu) = \mathbb{V}(\ell,\nu)$. Mappings $\mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$ can be compared, by using a pointwise comparison: in particular the maximum or minimum of two such mappings is defined pointwisely. The previous operators indeed allow us to split the computation of $\mathcal{F}_p$. We also rely on the classical notion of *regions*, as introduced in the seminal work on timed automata [2]. Indeed, for a given location ℓ, the set of deadlock valuations ν where $E(\ell,\nu) = \emptyset$ is a union of regions that we denote R_ℓ in the following, and that can easily be computed.

▶ **Lemma 11.** *For all $\mathbb{V}\colon L \times \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$, $\ell \in L$, and $p > 0$, $\mathcal{F}_p(\mathbb{V})(\ell)$ equals*

$$\begin{cases} \nu \mapsto 0 & \text{if } \ell \in L_T \\ \left(\nu \mapsto \begin{cases} +\infty & \text{if } \nu \in R_\ell \\ \left(\max\limits_{\delta=(\ell,g,Y,\ell')\in\Delta} \left[\mathsf{wt}(\delta) + \mathsf{Pre}_\ell(\mathsf{Guard}_\delta(\mathsf{Unreset}_Y(\mathbb{V}_{\ell'})))\right]\right)(\nu) & \text{if } \nu \notin R_\ell \end{cases}\right) & \text{if } \ell \in L_{\mathsf{Max}} \\ \min\limits_{\delta=(\ell,g,Y,\ell')\in\Delta} \left[\mathsf{wt}(\delta) + \mathsf{Pre}_\ell(\mathsf{Perturb}^p_\ell(\mathsf{Guard}_\delta(\mathsf{Unreset}_Y(\mathbb{V}_{\ell'}))))\right] & \text{if } \ell \in L_{\mathsf{Min}} \end{cases}$$

5 Encoding parametric piecewise affine functions

We now explain how to encode the mappings that the operators defined in the previous section take as input, to compute $\mathcal{F}_p$ for all perturbation bounds $p > 0$ at once. We adapt the formalism used in [1, 18] to incorporate the perturbation p. This formalism relies on the remark that $\mathbb{V}^0$ is a piecewise affine function, and that if $\mathbb{V}$ is piecewise affine, so is $\mathcal{F}_p(\mathbb{V})$: thus we only have to manipulate such piecewise affine functions.

To model the robustness, that depends on the perturbation bound p, and maintain a parametric description of all the value functions for infinitesimally small values of p, we consider piecewise affine functions that depend on a formal parameter $\mathfrak{p}$ describing the perturbation. The pieces over which the function is affine, that we call cells in the following, are polytopes described by a conjunction of affine equalities and inequalities involving $\mathfrak{p}$. Some of our computations will only hold for small enough values of the parameter $\mathfrak{p}$ and we will thus also maintain an upperbound for this parameter.

▶ **Definition 12.** *We call* parametric affine expression *an expression E of the form $\sum_{x\in\mathcal{X}} \alpha_x\, x + \beta + \gamma\mathfrak{p}$ with $\alpha_x \in \mathbb{Q}$ for all $x \in \mathcal{X}$, $\beta \in \mathbb{Q} \cup \{-\infty, +\infty\}$, and $\gamma \in \mathbb{Q}$. The semantics of such an expression is given for a particular perturbation p as a mapping $[\![E]\!]_p\colon \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$ defined for all $\nu \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$ by $[\![E]\!]_p(\nu) = \sum_{x\in\mathcal{X}} \alpha_x\, \nu(x) + \beta + \gamma p$.*

A partition of $\mathbb{R}^{\mathcal{X}}_{\geq 0}$ into cells is described by a set $\mathcal{E} = \{E_1, \ldots, E_m\}$ of parametric affine expressions. Every expression can be turned into an equation or inequation by comparing it to 0 with the symbol $=$, $<$ or $>$. The partition of $\mathbb{R}^{\mathcal{X}}_{\geq 0}$ is obtained by considering all

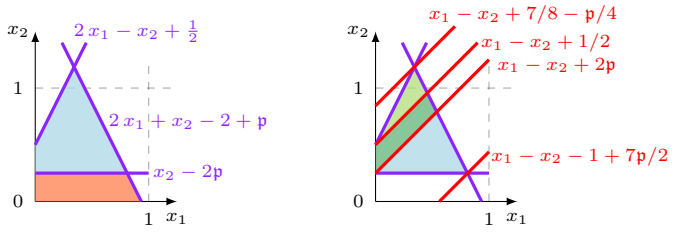

Figure 3 On the left, we depict the partition defined from $\mathcal{E} = \{x_2 - 2\mathfrak{p}, 2x_1 + x_2 - 2 + \mathfrak{p}, 2x_1 - x_2 + 1/2\}$, for a small enough value of $\mathfrak{p}$. On the right, we depict the atomic partition induced by $\mathcal{E}$, and draw in red the added parametric affine expressions.

the combinations of equations and inequations for each $1 \leq i \leq m$: such a combination is described by a tuple $(\bowtie_i)_{1\leq i\leq m}$ of symbols in $\{=, <, >\}$. For a given perturbation p, we let $[\![\mathcal{E}, (\bowtie_i)_{1\leq i\leq m}]\!]_p$ be the set of valuations ν such that for all $i \in \{1, \ldots, m\}$, $[\![E_i]\!]_p \bowtie_i 0$.

We call *cell* every such combination such that for p that tends to 0, while being positive, the set $[\![\mathcal{E}, (\bowtie_i)_{1\leq i\leq m}]\!]_p$ is non empty. We let $\mathcal{C}(\mathcal{E})$ be the set of cells of $\mathcal{E}$. Notice that it can be decided (in at most exponential time) if a combination $(\bowtie_i)_{1\leq i\leq m}$ is a cell, by encoding the semantics in the first order theory of the reals, and deciding if there exists an upperbound $\eta > 0$ such that for all $0 < p \leq \eta$, $[\![\mathcal{E}, (\bowtie_i)_{1\leq i\leq m}]\!]_p$ is non empty. Moreover, we can compute the biggest such upperbound η if it exists. The upperbound of the partition $\mathcal{E} = \{E_1, \ldots, E_m\}$ is then defined as the minimum such upperbound over all cells (there are at most 3^m cells), and denoted by $\eta(\mathcal{E})$ in the following. On the left of Figure 3, we depict the partition of $\mathbb{R}^{\mathcal{X}}_{\geq 0}$ defined from $\mathcal{E} = \{x_2 - 2\mathfrak{p}, 2x_1 + x_2 - 2 + \mathfrak{p}, 2x_1 - x_2 + 1/2\}$, with a fixed value of the perturbation. In blue, we color the cell defined by $(>, <, >)$. We note that this cell is non-empty when $p \leq 1/2$. By considering other cells, we obtain that $\eta(\mathcal{E}) = 1/2$.

In the following, we may need to record a smaller upperbound than $\eta(\mathcal{E})$, in order to keep the tightest constraint seen so far in the computation. We thus call *parametric partition* a pair $\langle \mathcal{E}, \eta \rangle$ given by a set of equations and a perturbation $\eta > 0$ that is at most $\eta(\mathcal{E})$.

For a cell $c \in \mathcal{C}(\mathcal{E})$, an expression E of $\mathcal{E}$ is said to be on the *border* of c if the removal of E from the set $\mathcal{E}$ of expressions forbids one to obtain the set of valuations $[\![c]\!]_p$ with the resulting cells for all small enough values of $p > 0$: more precisely, we require that no cell $c' \in \mathcal{C}(\mathcal{E} \setminus \{E\})$ is such that for some $p \leq \eta(\mathcal{E})$, $[\![c]\!]_p = [\![c']\!]_p$. Because of the definition of $\eta(\mathcal{E})$, this definition does not depend on the actual value of p that we consider (and we could thus replace "for some" by "for all" above). On the left of Figure 3, all expressions are on the border for the blue cell, but only two of them are on the border of the orange cell.

The proofs that follow (in particular time delaying that requires to move along diagonal lines) requires to adapt the notion of *atomicity* of a parametric partition, originally introduced in the non-robust setting [1, 18]. A parametric affine expression $E = \sum_{x\in\mathcal{X}} \alpha_x\, x + \beta + \gamma p$ is said to be *diagonal* if $\sum_{x\in\mathcal{X}} \alpha_x = 0$: indeed, for all $p > 0$, $[\![E]\!]_p(\nu) = [\![E]\!]_p(\nu + t)$ for all $t \in \mathbb{R}$. A parametric partition is said to be *atomic* if for all cells $c \in \mathcal{C}(\mathcal{E})$, there are at most two non-diagonal parametric affine expressions on the border of c: intuitively, one border is reachable from every valuation by letting time elapse, and the other border is such that by letting time elapse from it we can reach all valuations of the cell. An atomic partition decomposes the space into tubes whose borders are only diagonal, each tube being then sliced by using only non-diagonal expressions. In particular, each cell c of an atomic partition has a finite set of cells that it can reach by time elapsing (and dually a set of cells that can reach c by time elapsing), and this set does not depend on the value of the parameter p, nor the

starting valuation in $[\![c]\!]_p$. On the right of Figure 3, we depict the atomic partition associated with the same set of expressions used on the left. The diagonal expressions are depicted in red. We note that the cell colored in blue on the left is split into five cells that are non-empty when $p \leq 3/7$. We can describe the new parametric partition as $(\{x_2 - 2\mathfrak{p}, 2x_1 + x_2 - 2 + \mathfrak{p}, 2x_1 - x_2 + 1/2, x_1 - x_2 - 1 + 7\mathfrak{p}/2, x_1 - x_2 + 2\mathfrak{p}, x_1 - x_2 + 1/2, x_1 - x_2 + 7/8 - \mathfrak{p}/4\}, 3/7)$. As we will see below, a parametric partition can always be made atomic, by adding some diagonal parametric affine expressions.

A *parametric value function* (PVF for short) is a tuple $F = \langle \mathcal{E}, \eta, (f_c)_{c \in \mathcal{C}(\mathcal{E})} \rangle$ where $\langle \mathcal{E}, \eta \rangle$ is a partition and, for all cells $c \in \mathcal{C}(\mathcal{E})$, f_c is a parametric affine expression. For a perturbation $0 < p \leq \eta$, the semantics $[\![F]\!]_p$ of this tuple is a mapping $\mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$ defined for all valuations ν by $[\![F]\!]_p(\nu) = [\![f_c]\!]_p(\nu)$ where c is the unique cell such that $\nu \in [\![c]\!]_p$. A PVF is said to be *atomic* if its parametric partition is atomic. As announced above, we can always refine a PVF so that it becomes atomic.

▶ **Lemma 13.** *If $F = \langle \mathcal{E}, \eta, (f_c)_{c \in \mathcal{C}(\mathcal{E})} \rangle$ is a PVF, we can compute an atomic PVF $F' = \langle \mathcal{E}', \eta', (f'_c)_{c \in \mathcal{C}(\mathcal{E}')} \rangle$ such that $\eta' \leq \eta$, and $[\![F]\!]_p = [\![F']\!]_p$ for all $p \leq \eta'$.*

To conclude the proof of Theorem 9, we need to compute one application of $\mathcal{F}_p$ over a mapping $X \colon L \times \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathbb{R}_\infty$ that is stored by a PVF for each location. Moreover, our computations must be done for all small enough values of p simultaneously.

▶ **Proposition 14.** *Let $F = \langle \mathcal{E}, \eta, (f_c)_{c \in \mathcal{C}(\mathcal{E})} \rangle$ be a PVF. We can compute a PVF $F' = \langle \mathcal{E}', \eta', (f'_c)_{c \in \mathcal{C}(\mathcal{E}')} \rangle$ with $\eta' \leq \eta$, and $[\![F']\!]_p = \mathcal{F}_p([\![F]\!]_p)$ for all $p \leq \eta'$.*

Sketch of the proof. By using Lemma 11, it suffices to perform the proof independently for the four operators, as well as maximum or minimum operations. Proofs from [1, 18] can be adapted for the maximum/minimum operations as well as the operators Guard_δ and $\mathsf{Unreset}_Y$ that exist also in the non-robust setting. In the case of Max, the two cases depend only on the regions and we thus only apply the various operators for starting valuations not in R_ℓ.

For the operator Pre_ℓ (and similarly $\mathsf{Perturb}^p_\ell$), the adaptation is more subtle. The key ingredient, for instance to compute it over an atomic partition for a location ℓ of Max, is to transform the computation of $([\![F]\!]_p)(\nu) = \sup_{t \geq 0}[t\,\mathsf{wt}(\ell) + [\![F]\!]_p(\nu + t)]$ involving a supremum (for a fixed valuation ν, and a fixed perturbation $p \leq \eta$), by using a maximum over a finite set of interesting delays. First, we remark that for every delay $t > 0$, the valuation $\nu + t$ belongs to the open diagonal half-line from ν, which crosses some of the semantics $[\![c']\!]_p$ for certain cells c'. Moreover, this finite (since there are anyway only a finite number of cells in the partition) subset of crossing cells neither depends on the choice of ν in a given starting cell c, nor on the perturbation p as long as it is at most η (by atomicity of the partition). Since the function $[\![F]\!]_p$ is affine in each cell, the above supremum over the possible delays is obtained for a value t that is either 0, or tending to $+\infty$, or on one of the two non-diagonal borders of the previous crossing cells, and we thus only have to consider those borders (that do not depend on the choice of ν in a given starting cell c, nor on the perturbation p). ◀

6 Divergent weighted timed games

From our algorithm to solve acyclic WTGs, we naturally want to extend the computation of the robust value to other classes of WTGs by using an unfolding of the WTG. In particular, we consider the natural extension of *divergent WTGs* (like in [16, 18]) that define a large class of decidable WTGs for the exact semantics, with no limitations on the number of clocks.

As usual in related work [1, 6, 7, 18], we now assume that all clocks are *bounded* by a constant $\mathsf{M} \in \mathbb{N}$, i.e. every transition of the WTG is equipped with a guard g such that $\nu \models g$ implies $\nu(x) \leq \mathsf{M}$ for all clocks $x \in \mathcal{X}$. We denote by W_{loc} (resp. W_{tr}, W_{e}) the maximal weight in absolute values of locations (resp. of transitions, edges) of $\mathcal{G}$, i.e. $W_{\mathsf{loc}} = \max_{\ell \in L} |\mathsf{wt}(\ell)|$ (resp. $W_{\mathsf{tr}} = \max_{\delta \in \Delta} |\mathsf{wt}(\delta)|$, $W_{\mathsf{e}} = \mathsf{M}W_{\mathsf{loc}} + W_{\mathsf{tr}}$).

We use the exact semantics to define the divergence property by relying once again on the regions [2]. We let $\mathsf{Reg}(\mathcal{X}, \mathsf{M})$ be the set of regions when clocks are bounded by M. A game $\mathcal{G}$ (w.r.t. the exact semantics) can be populated with the region information as described formally in [18]: we obtain the *region game* $\mathcal{R}(\mathcal{G})$. We call *region path* a finite or infinite sequence of transitions in this game, and we denote by π such paths. A play ρ in $\mathcal{G}$ can be projected on a region path π: we say that ρ *follows* the region path π. It is important to notice that, even if π is a *cycle* (i.e. starts and ends in the same location of the region game), there may exist plays following it in $\mathcal{G}$ that are not cycles, due to the fact that regions are sets of valuations.

Divergent WTGs are obtained by enforcing a semantical property of divergence (originally called *strictly non-Zeno cost* when only dealing with non-negative weights [6]): it asks that every play (w.r.t. the exact semantics) following a cyclic region path has weight far from 0. Formally, a cyclic region path π of $\mathcal{R}(\mathcal{G})$ is said to be a positive (resp. negative) if every finite play ρ following π satisfies $\mathsf{wt}_\Sigma(\rho) \geq 1$ (resp. $\mathsf{wt}_\Sigma(\rho) \leq -1$).

▶ **Definition 15.** *A WTG is divergent if every cyclic region path is positive or negative.*

Finally, with loss of generality, we only consider divergent WTGs containing no configurations with a value equal to $-\infty$. Intuitively, guaranteeing a value $-\infty$ resembles a Büchi condition for Min, since this means that Max cannot avoid the iteration of negative cycles with his delays. In the robust settings, testing Büchi condition for automata is already non-trivial [19], thus we remove this behaviour in our games in this article. Since the value is a lower bound of the robust value (by Lemma 7), we obtain that all locations have a robust value distinct from $-\infty$. Moreover, testing if such a location exists in a divergent WTG can be done in $\mathsf{EXPTIME}$ [18]. Our second contribution is to extend the symbolic algorithm used in the case of acyclic WTGs to compute the robust value in this subclass of divergent WTGs.

▶ **Theorem 16.** *The robust value problem is decidable over the subclass of divergent WTGs without configurations of value $-\infty$.*

Sketch of the proof. We compute the robust value by using an adaptation of the algorithm of [18] used to compute the (exact) value function in divergent WTGs. In particular, its termination is guaranteed by the use of an equivalent definition of divergent WTG requiring that for all strongly connected components (SCC) S of the graph of the region game, either every cyclic region path π inside S is positive (we say that the SCC is positive) or every cyclic region path π inside S is negative (we say that the SCC is negative).

We adapt this argument in the case of the computation of the robust value of a divergent WTG (without configurations with a value equal to $-\infty$). In particular, we observe that if a cyclic region path is positive (resp. negative) w.r.t. the exact semantics, then it is also positive (resp. negative) w.r.t. the conservative semantics, as the latter only filters some plays. Thus, the finite convergence of the value iteration algorithm (defined by $\mathcal{F}_p$ as for acyclic WTG, i.e. initialised by the function $\mathbb{V}^0$ defined such that $\mathbb{V}^0(\ell) = 0$ for all target locations, and $\mathbb{V}^0(\ell) = +\infty$ otherwise) is guaranteed by its finite convergence in finite time in each positive (resp. negative) SCC. Intuitively, in a positive (resp. negative) SCC, the interest of Min (resp. Max) is to quickly reach a target location of $\mathcal{G}$ to minimise the number of positive

(resp. negative) cyclic region paths followed along the play that allow us to upperbound the number of iterations needed to obtain the robust value of all locations. Thus, the number of iterations needed to compute the robust value in a divergent WTG is defined by the sum of the number of iterations for each SCC along the longest path of the SCC decomposition. ◀

On top of computing the value, the modified algorithm allows one to synthesize almost-optimal strategies (we can adapt recent works [23] showing that those strategies can be taken among switching strategies for Min and memoryless strategies for Max).

7 Conclusion

This article allows one to compute (finite) robust values of weighted timed games in classes of games (acyclic and divergent) where the non-robust values are indeed computable.

As future works, we would like to carefully explore the exact complexity of our algorithms. Intuitively, each operator used to describe $\mathcal{F}_p$ can be computed in exponential time with respect to the set of cells and the size of η. By [18], the number of cells exponentially grows at each application of $\mathcal{F}_p$ (so it is doubly exponential for the whole computation) and the constants in affine expressions polynomially grow, in the non-robust setting. We hope that such upperbounds remain in the robust setting. This would imply that, for divergent WTGs, our algorithm requires a triply-exponential time, since the unfolding is exponential in the size of $\mathcal{G}$.

We also suggest to extend the setting to incorporate divergent WTGs that contain location with a value equal to $-\infty$. However, fixing it for all divergent WTGs seems to be difficult since, intuitively, the condition to guarantee $-\infty$ looks like a Büchi condition where Max can avoid the iteration of cycles with his delays. Moreover, we would like to study almost-divergent weighted timed games (studied in [17, 18]), a class of games undecidable, but with approximable value functions. We wonder if the robust values could also be approximated by similar techniques. Another direction of research would be to consider the fragment of one-clock weighted timed games, another class of games where the value function is known to be computable (for a long time in the non-negative case [5], very recently in the general case [24]). The difficulty here would be that encoding the conservative semantics in an exact semantics, even with fixed-perturbation, requires the addition of a clock, thus exiting the decidable fragment. The question thus becomes a possible adaptation of techniques used previously to solve non robust one-clock WTGs to incorporate directly the robustness.

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 *LNCS*, 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 Eugene Asarin and Oded Maler. As soon as possible: Time optimal control for timed automata. In *Hybrid Systems: Computation and Control*, volume 1569 of *LNCS*, pages 19–30. Springer, 1999. doi:10.1007/3-540-48983-5_6.

4 Patricia Bouyer, Thomas Brihaye, Véronique Bruyère, and Jean-François Raskin. On the optimal reachability problem of weighted timed automata. *Formal Methods in System Design*, 31(2):135–175, 2007.

5 Patricia Bouyer, Thomas Brihaye, and Nicolas Markey. Improved undecidability results on weighted timed automata. *Information Processing Letters*, 98(5):188–194, 2006.

6 Patricia Bouyer, Franck Cassez, Emmanuel Fleury, and Kim G. Larsen. Optimal strategies in priced timed game automata. In *FSTTCS 2004: Foundations of Software Technology and Theoretical Computer Science*, LNCS, pages 148–160, Berlin, Heidelberg, 2005. Springer. doi:10.1007/978-3-540-30538-5_13.

7 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. Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.CONCUR.2015.311.

8 Patricia Bouyer, Nicolas Markey, and Pierre-Alain Reynier. Robust model-checking of linear-time properties in timed automata. In *Proceedings of the 7th Latin American Conference on Theoretical Informatics (LATIN'06)*, LNCS, pages 238–249. Springer, 2006. doi:10.1007/11682462_25.

9 Patricia Bouyer, Nicolas Markey, and Ocan Sankur. Robust weighted timed automata and games. In *Proceedings of the 11th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'13)*, volume 8053 of *LNCS*, pages 31–46. Springer, August 2013. doi:10.1007/978-3-642-40229-6_3.

10 Patricia Bouyer, Nicolas Markey, and Ocan Sankur. Robust reachability in timed automata: Game-based approach. *Journal of Theoretical Computer Science (TCS)*, 563:43–74, 2015.

11 Thomas Brihaye, Véronique Bruyère, and Jean-François Raskin. On optimal timed strategies. In *Formal Modeling and Analysis of Timed Systems*, LNCS, pages 49–64. Springer, 2005. doi:10.1007/11603009_5.

12 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 Foundations 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.

13 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), August 2022. doi:10.46298/lmcs-18(3:17)2022.

14 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 *LNCS*, pages 560–575. Springer, 2014. doi:10.1007/978-3-662-44584-6_38.

15 Damien Busatto-Gaston. *Symbolic controller synthesis for timed systems: robustness and optimality.* PhD thesis, Aix-Marseille Université, 2019.

16 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 2017)*, LNCS, pages 162–178. Springer, 2017. doi:10.1007/978-3-662-54458-7_10.

17 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, December 2018. doi:10.4230/LIPIcs.FSTTCS.2018.28.

18 Damien Busatto-Gaston, Benjamin Monmege, and Pierre-Alain Reynier. Optimal controller synthesis for timed systems. *Log. Methods Comput. Sci.*, 19(1), 2023. doi:10.46298/LMCS-19(1:20)2023.

19 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 (CAV 2019)*, volume 11561 of *LNCS*, pages 572–590. Springer, 2019. doi:10.1007/978-3-030-25540-4_33.

20 Martin De Wulf, Laurent Doyen, Nicolas Markey, and Jean-François Raskin. Robustness and implementability of timed automata. In *Proceedings of the International Conference on Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems (FORMATS'2004)*, LNCS, pages 118–133. Springer, 2004. doi:10.1007/978-3-540-30206-3_10.

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

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

23 Benjamin Monmege, Julie Parreaux, and Pierre-Alain Reynier. Playing Stochastically in Weighted Timed Games to Emulate Memory. In *Proceedings of the 48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, volume 198 of *LIPIcs*, pages 137:1–137:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICALP.2021.137.

24 Benjamin Monmege, Julie Parreaux, and Pierre-Alain Reynier. Decidability of one-clock weighted timed games with arbitrary weights. In *Proceedings of the 33rd International Conference on Concurrency Theory (CONCUR 2022)*, volume 243 of *LIPIcs*, pages 15:1–15:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICS.CONCUR.2022.15.

25 Benjamin Monmege, Julie Parreaux, and Pierre-Alain Reynier. Synthesis of robust optimal strategies in weighted timed games. *CoRR*, abs/2403.06921, 2024. doi:10.48550/arXiv.2403.06921.

26 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 *LNCS*, pages 203–217. Springer, 2014. doi:10.1007/978-3-662-44584-6_15.

27 Anuj Puri. Dynamical properties of timed automata. *Discrete Event Dynamic Systems*, 10:87–113, 2000. doi:10.1023/A:1008387132377.

28 Ocan Sankur. *Robustness in timed automata : analysis, synthesis, implementation. (Robustesse dans les automates temporisés : analyse, synthèse, implémentation)*. PhD thesis, École normale supérieure de Cachan, Paris, France, 2013.

29 Ocan Sankur, Patricia Bouyer, and Nicolas Markey. Shrinking Timed Automata. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2011)*, volume 13 of *LIPIcs*, pages 90–102. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2011. doi:10.4230/LIPIcs.FSTTCS.2011.90.

30 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 2013)*, volume 8052 of *LNCS*, pages 546–560. Springer, 2013. doi:10.1007/978-3-642-40184-8_38.