



HAL
open science

Optimal controller synthesis for timed systems

Damien Busatto-Gaston, Benjamin Monmege, Pierre-Alain Reynier

► **To cite this version:**

Damien Busatto-Gaston, Benjamin Monmege, Pierre-Alain Reynier. Optimal controller synthesis for timed systems. Logical Methods in Computer Science, 2023, Volume 19, Issue 1, 10.46298/lmcs-19(1:20)2023 . hal-04442588

HAL Id: hal-04442588

<https://hal.science/hal-04442588>

Submitted on 6 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

OPTIMAL CONTROLLER SYNTHESIS FOR TIMED SYSTEMS

DAMIEN BUSATTO-GASTON ^a, BENJAMIN MONMEGE ^b, AND PIERRE-ALAIN REYNIER ^b

^a Univ Paris Est Creteil, LACL, F-94010 Creteil, France
e-mail address: damien.busatto-gaston@u-pec.fr

^b Aix Marseille Univ, LIS, CNRS, Marseille, France
e-mail address: benjamin.monmege@univ-amu.fr, pierre-alain.reynier@univ-amu.fr

ABSTRACT. Weighted timed games are zero-sum games played by two players on a timed automaton equipped with weights, where one player wants to minimise the cumulative weight while reaching a target. Used in a reactive synthesis perspective, this quantitative extension of timed games allows one to measure the quality of controllers in real-time systems. Weighted timed games are notoriously difficult and quickly undecidable, even when restricted to non-negative weights. For non-negative weights, a fragment of weighted timed games that can be analysed has been introduced by Bouyer, Jaziri and Markey in 2015. Though the value problem is undecidable even in this fragment, the authors show how to approximate the value by considering regions with a refined granularity. In this work, we extend this class to incorporate negative weights, allowing one to model energy for instance, and prove that the value can still be approximated, with the same complexity (provided that clocks are bounded). A restriction also allows us to obtain a class of decidable weighted timed games with negative weights and an arbitrary number of clocks. In addition, we show that a symbolic algorithm, relying on the paradigm of value iteration, can be used as an approximation/computation schema over these classes. We also consider the special case of untimed weighted games, where the same fragments are solvable in polynomial time: this contrasts with the pseudo-polynomial complexity, known so far, for weighted games without restrictions.

1. INTRODUCTION

We are interested in the design of programs sensitive to real-time, where keeping track of how much time elapses between the decisions taken by the program is required to differentiate the good and bad behaviours of the system. This is a common requirement for embedded systems, as they interact with the real world. The design of such programs is a notoriously difficult problem, because they must take care of delicate timing issues, and are difficult to debug *a posteriori*. In order to ease the design of real-time software, it appears important to automatise the process by using formal methods. The situation may be modelled into a

Key words and phrases: weighted timed games; computational game theory; timed automata; quantitative constraints.

We thank the reviewers of this article, as well as the two conference versions from which most of the results come from, for their valuable feedback.

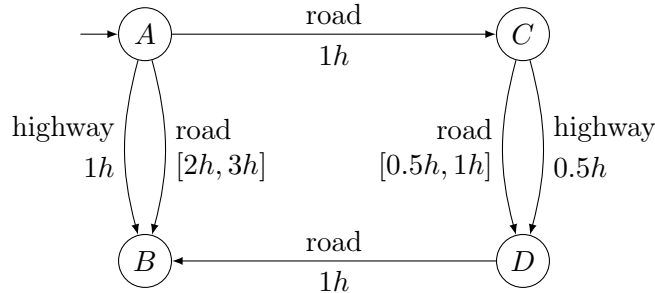


Figure 1: A ride-sharing decision diagram.

timed automaton [AD94], namely a transition system equipped with real-valued variables, called clocks, evolving with a uniform rate. Transitions are equipped with timing constraints expressed over the clocks, and may only be taken when these constraints are met. Finally, clocks may be reset so as to keep track of how much time has elapsed since a particular transition was taken.

In order to check whether the real-time system fulfils a certain specification, one determines whether there exists an accepting execution in the timed automaton: this is the classical *model-checking problem*. A simple, yet realistic specification asks that a target state is reached at some point. It has been proven in early works [AD94] that the reachability problem on timed automata is PSPACE-complete: in particular, the PSPACE upper bound is obtained by partitioning the state space into a finite number of *regions*. While optimal from a theoretical complexity point of view, practical tools tend to favour efficient *symbolic* algorithms for solving such model-checking problems like reachability, that use *zones* instead of regions, as they allow an on-demand partitioning of the state space. This leads to much better performances, as witnessed by successful model-checking tools like UPPAAL [LPY97], KRONOS [BDM⁺98], or TCHECKER [HPT19, HSW10].

Instead of verifying a system, we can try to synthesise one automatically. One setting consists in using *game theory*. The set of configurations of the system is then partitioned into two *players*: a controller whose role is to fulfil the specification, and an antagonistic environment. The goal becomes to find automatically a good *strategy* for the controller, which is called the *controller synthesis problem*. In timed systems, both players alternatively choose transitions and delays in a timed automaton: this is called a *timed game*. Strategies of players are recipes dictating how to play (timing delays and transitions to follow). In this ambitious setting, we will focus on reachability objectives, and we are thus looking for a strategy of the controller so that the target is reached no matter how the environment plays. Reachability timed games are decidable [AM99], and EXPTIME-complete [JT07].

If the controller has a winning strategy in a given reachability timed game, several such winning strategies could exist. Weighted extensions of these games have been considered in order to measure the quality of the winning strategy for the controller [BCFL04]. This means that the game now takes place over a *weighted (or priced) timed automaton* [BFH⁺01, ALTP04], where edges are equipped with weights, and locations with rates of weights (the cost is then proportional to the time spent in this location, with the rate as proportional coefficient). The optimal reachability problem asks what is the lowest cumulative weight that the controller can guarantee for reaching a target from a given initial state, against any decision made by the antagonistic environment. The controller is therefore the minimiser

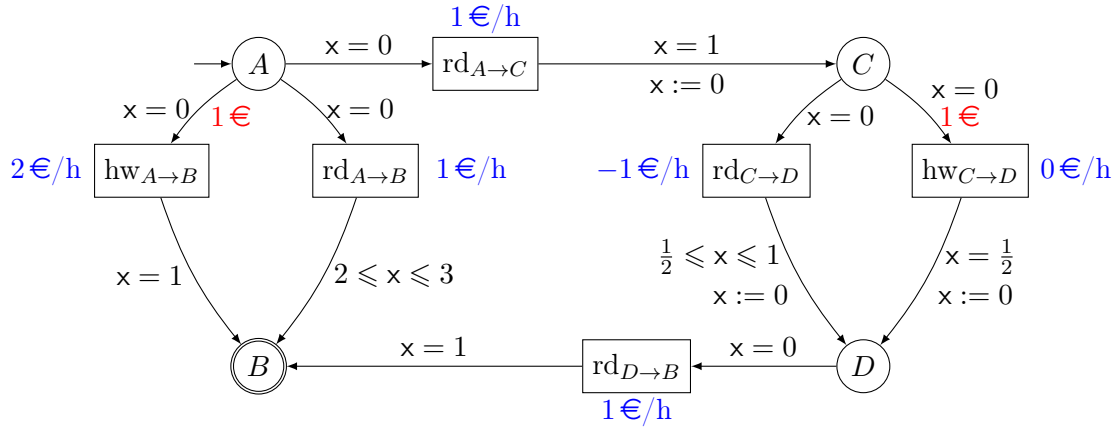


Figure 2: A weighted timed game modelling Figure 1. Transitions are labelled by a guard over clock x and by the reset $x := 0$ when needed. The cost of waiting in a state is displayed in blue, the cost of taking a transition is in red. States and transitions without costs have a weight of 0€

player, while the environment wants to maximise the weight accumulated along the execution. This lowest weight is called the *value* of the game, and computing it can be seen as a natural generalisation of the classical shortest path problem in a weighted graph to the case of two-player timed games.

Example 1.1. As a motivating example, we present a ride-sharing scenario. As a driver, we wish to travel from point A to point B , and must choose between several options, as displayed in Figure 1. We can use a direct road, and reach B in two to three hours, or a highway that lets us reach our destination in one hour. Alternatively, we can make a detour: another traveller is waiting at point C , and wishes to reach point D . For this portion too, a faster highway is available.

While all four possible paths satisfy the objective "reaching B ", we want to select the one that lets us spend as little money as possible for the trip. The cost of each path depends on several factors. There are fixed entry fees (of 1€) for the highways, and we need to keep track of fuel consumption, as the rate at which fuel is used differs in roads and highways. Thus, we say that roads cost 1€ per hour, while highways cost 2€ per hour. Moreover, if we share the portion from C to D , the other traveller will pay for his trip (at a rate of $2 \text{€}/h$), and that can lower our costs. A shared road therefore costs us -1€ per hour (negative rate means we are making a profit), while a shared highway costs $0 \text{€}/h$.

The situation can be modelled as a weighted timed game, displayed in Figure 2. The controller chooses delays and transitions in circle states, while the environment controls the square ones. For example, if we choose to use the direct road from A to B , we go (immediately) to state $rd_{A \rightarrow B}$, and stay there until going to state B . This requires letting between two and three hours elapse in $rd_{A \rightarrow B}$, with a cost of $1 \text{€}/h$. The delay is chosen by the environment, as it depends on external influences like traffic density.

In this example, the optimal strategy is to share the road from C to D . This lets us ensure a total weight of at most 1.5€ : going from A to C costs 1€ in the worst case; going

from D to B similarly costs at most 1€ ; and sharing the trip from C to D is guaranteed to bring us at least 0.5€ .

While deciding whether the value of a weighted timed automata is lower than a given threshold has been shown to be PSPACE-complete [BBBR07] (*i.e.* the same complexity as the non-weighted version, that is the reachability problem), the same problem is known to be undecidable in weighted timed games [BBR05]. This justifies the study of restrictions in order to regain decidability, the first and most interesting one being the class of strictly non-Zeno cost with only non-negative weights (in edges and locations) [BCFL04]: this hypothesis states that every execution of the timed automaton that follows a cycle of the region abstraction has a weight far from 0 (in the interval $[1, +\infty)$, for instance).

Less is known for weighted timed games in the presence of negative weights in edges and locations. In particular, no results existed before this work for a class that does not restrict the number of clocks of the timed automaton to 1. However, several clocks are needed to keep track of how much time elapsed since multiple reference points, and negative weights are particularly interesting from a modelling perspective, for instance in case weights represent the consumption level of a resource (money, energy. . .) with the possibility to spend and gain some resource. We thus introduce a generalisation of the strictly non-Zeno cost hypothesis in the presence of negative weights, that we call *divergence*. Under the hypothesis that clocks are bounded, we show the decidability of the class of divergent weighted timed games for the optimal synthesis problem. The technique uses a *value iteration* procedure to solve weighted timed games for a bounded horizon, *i.e.* when controller has a fixed number of steps to reach his targets. It follows closely the framework of [ABM04], but is more symbolic and allows for negative weights. We then show that optimal strategies in divergent weighted timed games can be restricted to a bounded horizon, that matches the one obtained in the non-negative case from the study of [BCFL04].

The techniques providing these decidability results cannot be extended if the conditions are slightly relaxed. For instance, if we add the possibility for an execution of the timed automaton following a cycle of the region automaton to have weight *exactly 0*, the decision problem is known to be undecidable [BJM15], even with non-negative weights only. For this extension, in the presence of non-negative weights only, it has been proposed an approximation schema to compute arbitrarily close estimates of the optimal weight that the controller can guarantee [BJM15]. To this end, the authors consider regions with a refined granularity so as to control the precision of the approximation.

Our contribution on the approximation front is two-fold. We extend the class considered in [BJM15] to the presence of negative weights, and provide an approximation schema for the resulting class of *almost-divergent games* (again under the hypothesis that clocks are bounded). We then show that the approximation can be obtained using a symbolic computation, that avoids an *a priori* refinement of regions. Table 1 summarises our results on weighted timed games.

The classes of weighted timed games that we study induce interesting classes of *finite weighted games* when there are no clocks, that can be solved with a lower (polynomial) complexity than arbitrary weighted games, see Table 2.

Other types of payoffs than the cumulative weight we study (*i.e.* total payoff) have been considered for weighted timed games. For instance, energy and mean-payoff timed games have been introduced in [BCR14]. They are also undecidable in general. Interestingly, a subclass called *robust timed games*, not far from our divergence hypothesis, admits decidability results for other payoffs. A weighted timed game is robust if, to say short, every simple cycle (cycle

Timed	weights in \mathbb{N}		
	divergent	almost-divergent	all WTG
Value pb.	2-EXPTIME [BCFL04]+[ABM04]	undecidable [BJM15]	undecidable [BBR05]
Approx. pb.	/	2-EXPTIME [BJM15]+[ABM04]	?
Value $+\infty$	EXPTIME-complete [BCFL04]+[JT07]		

Timed	weights in \mathbb{Z}		
	divergent	almost-divergent	all WTG
Value pb.	3-EXPTIME EXPTIME-hard Thm. 3.4	undecidable [BJM15]	undecidable [BGNK ⁺ 14]
Approx. pb.	/	3-EXPTIME Thm. 3.9	?
Value $-\infty$	EXPTIME-complete Prop. 3.11,5.5		undecidable Prop. 3.12
Value $+\infty$	EXPTIME-complete Prop. 5.3		
Membership	PSPACE-complete Thm. 3.4,3.9		/

Table 1: Solving weighted timed games with arbitrary weights

Untimed	weights in \mathbb{N}	weights in \mathbb{Z}		
	all games	divergent	almost-div.	all games
Value pb.	PTIME-complete [KBB ⁺ 08], Section 11.2.2	PTIME-complete Thm. 11.4,11.6		pseudo-poly. [BGHM16]
Value $-\infty$	/	PTIME-complete Section 11.2.1		pseudo-poly. [BGHM16]
Value $+\infty$	PTIME-complete [KBB ⁺ 08], Section 11.2.2	PTIME-complete Prop. 5.3, [BGHM16]		
Membership	/	NL-complete (unary wt.) PTIME (binary wt.) Thm. 11.4,11.6		/

Table 2: Solving finite weighted games with arbitrary weights

without repetition of a state) has a weight that is either non-negative or less than $-\varepsilon$, for a fixed constant $\varepsilon > 0$ (the same constant for all cycles). Solving robust timed game can be done in EXPSpace, and is EXPTIME-hard. Moreover, deciding if a weighted timed game is robust has complexity 2-EXPSpace (and coNEXPTIME-hard). This contrasts with our PSPACE results for the membership problem.¹ It has to be noted that extending our techniques and results to the case of robust timed games may not be possible: indeed, with weights in \mathbb{N}

¹While our divergent games have a similar definition, the two classes are incomparable.

every game is robust, making the value problem for this class undecidable [BBR05], with no known approximation method.

This article is structured as follows. After some preliminary definitions of the models we study in Section 2, we introduce the divergent and almost-divergent classes of weighted timed games and state our results in Section 3. We study structural properties of these classes in order to be able to decide their membership in Section 4. We then start solving these games by studying the qualitative problem of deciding infinite values, using the crucial and central notion of kernels, in Section 5. The approximation schema for almost-divergent timed games is presented in Section 6 where we introduce and prove the correctness of a semi-unfolding of the game. Section 7 explains how to compute the value of this semi-unfolding in the special case of a tree-shaped game (without kernels) as previously studied in [ABM04]. We then generalise our study in Section 8 to take kernels into account; we also consider the special case of divergent timed games. The first approximation/computation we propose requires to first compute strongly connected components and to build the semi-unfolding of the timed games, which might be highly prohibitive as usual in the timed setting. Therefore, we present a more symbolic algorithm in Section 9. In Section 10, we deduce from our study an algorithm to synthesise finite-memory almost-optimal strategies in divergent timed games. Finally, we treat the special case of almost-divergent untimed games in Section 11, where polynomial time algorithms are provided in this case (to be compared with the pseudo-polynomial complexity in general known so far).

This work is based on works published in [BGMR18, BGMR17]: compared to these preliminary works, this article contains full and corrected proofs, as well as a detailed study of the acyclic case in Section 7, used as a building block for our study. We have also extended the work in the untimed case (Section 11) to incorporate almost-divergent weighted games.

2. PRELIMINARIES

2.1. Modelling real-time constraints. We first introduce notions that let us express timing constraints, useful to then define weighted timed games, and introduce classical tools for their study.

Let $\mathcal{X} = \{x_1, \dots, x_n\}$ be a finite, non-empty set of variables called clocks. A *valuation* $\nu: \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ is a mapping from clocks to non-negative real numbers, such that $\nu(x_1), \dots, \nu(x_n)$ are called the coordinates of ν . Equivalently, ν can be seen as a point in space $\mathbb{R}_{\geq 0}^{\mathcal{X}}$. We denote $\mathbf{0}$ the valuation such that for all $x \in \mathcal{X}$, $\nu(x) = 0$. Given a real number $d \in \mathbb{R}$, we define $\nu + d$ as the valuation such that $\forall x \in \mathcal{X}$, $(\nu + d)(x) = \nu(x) + d$ if it exists.² If d is non-negative, we say that we performed a time elapse of delay d . The time-successors of ν are the valuations $\nu + d$ with $d \geq 0$. Similarly, we refer to all $\nu + d$ in $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ with $d \leq 0$ as time-predecessors of ν . The set of points that are either time-predecessors or time-successors of a valuation ν form the unique diagonal line in $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ that contains ν . If \mathcal{Y} is a subset of \mathcal{X} , we define $\nu[\mathcal{Y} := 0]$ as the valuation such that $\forall x \in \mathcal{Y}$, $(\nu[\mathcal{Y} := 0])(x) = 0$ and $\forall x \in \mathcal{X} \setminus \mathcal{Y}$, $(\nu[\mathcal{Y} := 0])(x) = \nu(x)$. This operation is called a reset of clocks \mathcal{Y} .

We extend those notions to sets of valuations in a natural way. The set of time-successors of $Z \subseteq \mathbb{R}_{\geq 0}^{\mathcal{X}}$, denoted $\text{PostTime}(Z)$, contains the valuations that are time-successors of valuations in Z . The reset of $Z \subseteq \mathbb{R}_{\geq 0}^{\mathcal{X}}$ by \mathcal{Y} , denoted $Z[\mathcal{Y} := 0]$, contains the valuations $\nu[\mathcal{Y} := 0]$ such that $\nu \in Z$.

²if d is negative, $\nu + d$ may not belong to $\mathbb{R}_{\geq 0}^{\mathcal{X}}$

The term *atomic constraint* will refer to an affine inequality in one of the following forms:

- A strict (respectively non-strict) *non-diagonal* atomic constraint over clock $x \in \mathcal{X}$ and constant $c \in \mathbb{Q}$ is an inequality of the form $x \bowtie c$ with $\bowtie \in \{>, <\}$ (respectively $\bowtie \in \{\geq, \leq\}$).
- A strict (respectively non-strict) *diagonal* atomic constraint over clocks x and $y \in \mathcal{X}$ and constant $c \in \mathbb{Q}$ is an inequality of the form $x - y \bowtie c$ with $\bowtie \in \{>, <\}$ (respectively $\bowtie \in \{\geq, \leq\}$).

Let \top and \perp denote two special atomic constraints, defined as $x \geq 0$ and $x < 0$ for an arbitrary $x \in \mathcal{X}$. A *guard* g over \mathcal{X} is a finite conjunction of atomic constraints over clocks in \mathcal{X} . In particular, guards let us define $x = c$ as shorthand for $x \leq c \wedge x \geq c$, and $c_1 < x < c_2$ as shorthand for $x > c_1 \wedge x < c_2$. A guard is said strict (respectively non-strict, diagonal, non-diagonal) if all of its atomic constraints are strict (respectively non-strict, diagonal, non-diagonal). $\text{Guards}(\mathcal{X})$ denotes the set of all guards over \mathcal{X} , and $\text{Guards}^{\text{nd}}(\mathcal{X})$ the subset of non-diagonal guards. For all constants $c \in \mathbb{Q}$ and $\bowtie \in \{\geq, \leq, >, <\}$, we say that valuation $\nu \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$ satisfies the atomic constraint $x \bowtie c$ (respectively $x - y \bowtie c$), and write $\nu \models x \bowtie c$ (respectively $\nu \models x - y \bowtie c$), if $\nu(x) \bowtie c$ (respectively $\nu(x) - \nu(y) \bowtie c$). We say that valuation $\nu \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$ satisfies guard g , and write $\nu \models g$, if ν satisfies all atomic constraints in g . For $g \in \text{Guards}(\mathcal{X})$, let $\llbracket g \rrbracket$ denote the set of all $\nu \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$ such that $\nu \models g$. Such sets are called *zones* and form convex polyhedra of $\mathbb{R}_{\geq 0}^{\mathcal{X}}$. A guard g is said satisfiable when the zone $\llbracket g \rrbracket$ is non-empty, and a zone is called rectangular when the associated guard is non-diagonal. The universal zone refers to $\llbracket \top \rrbracket = \mathbb{R}_{\geq 0}^{\mathcal{X}}$ and the empty zone refers to $\llbracket \perp \rrbracket = \emptyset$. Guard \bar{g} is the closed version of a satisfiable guard g where every strict constraint of comparison operator $<$ or $>$ is replaced by its non-strict version \leq or \geq . The zone $\llbracket \bar{g} \rrbracket$ is the topological closure of $Z = \llbracket g \rrbracket$, and is also denoted \bar{Z} .

We will restrict ourselves to *bounded* clocks, so that clocks valuations will be point in $[0, M)^{\mathcal{X}}$ instead of $\mathbb{R}_{\geq 0}^{\mathcal{X}}$, for some upper bound $M \in \mathbb{N}_{>0}$.³ We denote by $\text{Guards}(\mathcal{X}, M)$ (respectively $\text{Guards}^{\text{nd}}(\mathcal{X}, M)$) the set of guards (respectively the set of non-diagonal guards) over \mathcal{X} bounded by M , in the sense that $|c| \leq M$ for all constants $c \in \mathbb{Q}$ appearing in the atomic constraints of the guards.

2.2. Regions. We will rely on the crucial notion of regions, as introduced in the seminal work on timed automata [AD94]. Let $\mathbb{Q}_N = \{a/N \mid a \in \mathbb{Z}\}$ be the set of rational numbers of granularity $1/N$ for a fixed $N \in \mathbb{N}_{>0}$. Given a finite set of rational numbers $\mathcal{S} \subseteq \mathbb{Q}$, \mathcal{S} is said to be of granularity $1/N$ if $\mathcal{S} \subseteq \mathbb{Q}_N$. Such an N always exists, and one can find the smallest one by decomposing elements of \mathcal{S} as irreducible fractions c/c' with $c \in \mathbb{Z}$, $c' \in \mathbb{N}_{>0}$ and use the least common multiple of all c' as N . A guard g is said to be of granularity $1/N$ if all constants in the atomic constraints of g form a set of granularity $1/N$. A zone is of granularity $1/N$ if it can be described by a guard of granularity $1/N$. Let $\text{Guards}_N(\mathcal{X}, M)$ denote the set of guards over \mathcal{X} bounded by M and of granularity $1/N$, and let $\text{Guards}_N^{\text{nd}}(\mathcal{X}, M)$ denote the non-diagonal ones. Given a finite set of guards $G \subseteq \text{Guards}^{\text{nd}}(\mathcal{X}, M)$, we can find N such that $G \subseteq \text{Guards}_N^{\text{nd}}(\mathcal{X}, M)$, by denoting $\mathcal{S} \subseteq \mathbb{Q}$ the set of constants used in atomic constraints of G and using N the smallest integer such that \mathcal{S} is of granularity $1/N$. For all $a \in \mathbb{R}_{\geq 0}$, $\lfloor a \rfloor \in \mathbb{N}$ denotes the integral part of a , and $\text{fract}(a) \in [0, 1)$ its fractional part, such that $a = \lfloor a \rfloor + \text{fract}(a)$.

³This assumption will be discussed and formalised later on in Hypothesis 3.

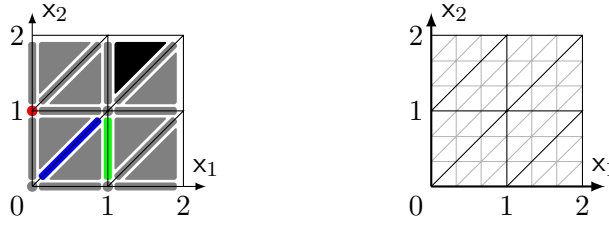


Figure 3: All $1/1$ -regions in $\text{Reg}_1(\{x_1, x_2\}, 2)$ on the left, their refinement of granularity $1/3$ in $\text{Reg}_3(\{x_1, x_2\}, 2)$ on the right.

Definition 2.1. With respect to the set \mathcal{X} of clocks, a granularity $N \in \mathbb{N}_{>0}$ and an upper bound $M \in \mathbb{N}_{>0}$ on the valuation of clocks, we partition the set $[0, M]^{\mathcal{X}}$ of valuations into $1/N$ -regions. We denote by $\text{Reg}_N(\mathcal{X}, M)$ the set of $1/N$ -regions bounded by M . Each such region is characterised by a pair (ι, β) where $\iota: \mathcal{X} \rightarrow [0, M] \cap \mathbb{Q}_N$ and β is a partition of \mathcal{X} into subsets $\beta_0 \uplus \beta_1 \uplus \dots \uplus \beta_m$ (with $m \geq 0$), where β_0 can be empty but $\beta_i \neq \emptyset$ for $1 \leq i \leq m$. A valuation ν of $[0, M]^{\mathcal{X}}$ belongs to the region characterised by (ι, β) if

- for all $x \in \mathcal{X}$, $\iota(x)N = \lfloor \nu(x)N \rfloor$;
- for all $x \in \beta_0$, $\text{fract}(\nu(x)N) = 0$;
- for all $0 \leq i \leq m$, for all $x, y \in \beta_i$, $\text{fract}(\nu(x)N) = \text{fract}(\nu(y)N)$;
- for all $0 \leq i < j \leq m$, for all $x \in \beta_i$ and all $y \in \beta_j$, $\text{fract}(\nu(x)N) < \text{fract}(\nu(y)N)$.

With granularity $N = 1$, we recover the classical notion of regions from [AD94] (in the case of bounded clocks), and we omit N from previous notations about regions, such that $1/N$ -regions are simply called regions, and $\text{Reg}_N(\mathcal{X}, M)$ is denoted $\text{Reg}(\mathcal{X}, M)$. The set of valuations contained in a $1/N$ -region r characterised by (ι, β) with $\beta = \beta_0 \uplus \beta_1 \uplus \dots \uplus \beta_m$ can be described by the guard $g_0 \wedge g_1 \wedge \dots \wedge g_m$ with

$$g_0 = \bigwedge_{x \in \beta_0} (x = \iota(x)), \quad g_1 = \bigwedge_{x, y \in \beta_1} (0 < x - \iota(x) = x - \iota(y) < 1/N)$$

and for $i \in \{2, \dots, m\}$,

$$g_i = \bigwedge_{x, y \in \beta_i} (z - \iota(z) < x - \iota(x) = y - \iota(y) < 1/N)$$

where z is any clock of β_{i-1} . Therefore, every $1/N$ -region is a zone of granularity $1/N$.

$\text{Reg}_N(\mathcal{X}, M)$ indeed forms a finite partition of $[0, M]^{\mathcal{X}}$. We bound the number of regions in the next lemma, making use of the fact that we only consider bounded regions of $[0, M]^{\mathcal{X}}$.

Lemma 2.2. *The number of $1/N$ -regions is polynomial in MN and exponential in the number of clocks $n = |\mathcal{X}|$:*

$$|\text{Reg}_N(\mathcal{X}, M)| \leq n!(2MN)^n.$$

Proof. Since $\iota(x) \in \{0, \frac{1}{N}, \dots, \frac{M-1}{N}\}$ for each clock x , there are $(MN)^n$ possible functions ι . For each of them, the partition β can be described by a total ordering of clocks (there are $n!$ such orderings), and a mapping of each clock x to either $=$ or $>$, so that the fractional part of $\nu(x)N$ is either equal to the fractional part of $\nu(y)N$ where y is the clock immediately preceding x in the ordering, or strictly greater than it. If the first clock of the ordering is mapped to $=$, it belongs to β_0 , otherwise $\beta_0 = \emptyset$. Therefore there are at most $n!2^n$ possibilities for the partition β . \square

If ν is a valuation in $[0, M)^{\mathcal{X}}$, $[\nu]$ denotes the unique region that contains ν . Valuations in the same $1/N$ -region r satisfy the same guards in $\text{Guards}_N(\mathcal{X}, M)$: we denote by $r \models g$ the satisfaction of the guard g for all valuations of the region r . In fact, zones associated to guards in $\text{Guards}_N(\mathcal{X}, M)$ can be described as a finite union of regions in $\text{Reg}_N(\mathcal{X}, M)$. If r is a $1/N$ -region in $\text{Reg}_N(\mathcal{X}, M)$, then the time-successor valuations in $\text{PostTime}(r) \cap [0, M)^{\mathcal{X}}$ form a finite union of regions in $\text{Reg}_N(\mathcal{X}, M)$, and the reset $r[\mathcal{Y} := 0]$ of $\mathcal{Y} \subseteq \mathcal{X}$ is a region in $\text{Reg}_N(\mathcal{X}, M)$. A $1/N$ -region r' is said to be a time successor of the $1/N$ -region r if there exists $\nu \in r$, $\nu' \in r'$, and $d > 0$ such that $\nu' = \nu + d$.

Example 2.3. Figure 3 represents the 24 regions of granularity $N = 1$ with upper bound $M = 2$ over two clocks. The green region is characterised by $\iota = (1, 0)$ and the partition β into two sets $\beta_0 = \{x_1\}$ and $\beta_1 = \{x_2\}$: it can be encoded with the guard $0 = x_1 - 1 < x_2 < 1$ and thus is equal to the zone $\llbracket x_1 = 1 \wedge 0 < x_2 < 1 \rrbracket$. The red region is characterised by $\iota = (0, 1)$ and the partition β into a single set $\beta_0 = \{x_1, x_2\}$. The black region is characterised by $\iota = (1, 1)$ and the partition β into three sets $\beta_0 = \emptyset$, $\beta_1 = \{x_1\}$ and $\beta_2 = \{x_2\}$. The blue region is characterised by $\iota = (0, 0)$ and the partition β into two sets $\beta_0 = \emptyset$ and $\beta_1 = \{x_1, x_2\}$.

2.3. Piecewise affine functions. We will heavily rely on the class of *piecewise affine functions*, and a way to efficiently encode them, as developed in [ABM04]. These functions will enable us to describe the value of weighted timed games and will therefore be restricted to a domain $[0, M)^{\mathcal{X}}$, with values taken in $\mathbb{R}_\infty = \mathbb{R} \cup \{+\infty, -\infty\}$. Let n denote the number of clocks, such that $\mathcal{X} = \{x_1, \dots, x_n\}$. An *affine function* is a mapping $f : [0, M)^{\mathcal{X}} \rightarrow \mathbb{R}_\infty$ such that for all $\nu \in [0, M)^{\mathcal{X}}$,

$$f(\nu) = a_1 \cdot \nu(x_1) + \dots + a_n \cdot \nu(x_n) + b$$

with partial derivatives $a_i \in \mathbb{Q}$ for $1 \leq i \leq n$, and additive constant $b \in \mathbb{Q}$. In this case, we say that f is defined by the equation $y = a_1 x_1 + \dots + a_n x_n + b$ where the variable $y \notin \mathcal{X}$ refers to $f(x_1, \dots, x_n)$. We also consider infinite mappings $\nu \mapsto +\infty$ and $\nu \mapsto -\infty$ to be affine functions, defined with null partial derivatives and an infinite constant b .

Intuitively, we define a piecewise affine function as a partition of $[0, M)^{\mathcal{X}}$ into finitely many polyhedra, called cells, each equipped by an affine function. Formally, an affine inequality is an equation I of the form

$$a_1 x_1 + \dots + a_n x_n + b \prec 0$$

where $b \in \mathbb{Q}$ is the additive constant of I , $\prec \in \{<, \leq\}$ is its comparison operator, and for every $1 \leq i \leq n$, $a_i \in \mathbb{Q}$ is the i -th partial derivative of I . Similarly, an affine equality is an equation E of the form

$$a_1 x_1 + \dots + a_n x_n + b = 0.$$

We say that $\nu \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$ satisfies I (respectively E), and write $\nu \models I$ (respectively $\nu \models E$), if $a_1 \cdot \nu(x_1) + \dots + a_n \cdot \nu(x_n) + b \prec 0$ (respectively $= 0$) holds. In this case, $\llbracket I \rrbracket$ (respectively $\llbracket E \rrbracket$) refers to the set of valuations that satisfy I (respectively E). Equalities (respectively inequalities) are equivalent when they are satisfied by the same valuations. In particular, multiplying the additive constant b and all partial derivatives a_i by the same factor $N \in \mathbb{N}_{>0}$ gives an equivalent equality (respectively inequality), and we will therefore assume that they are always integers.

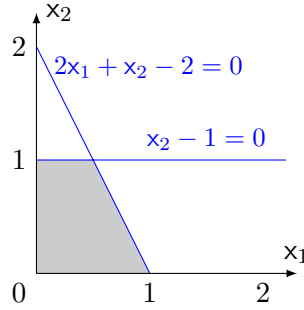


Figure 4: The cell $2x_1 + x_2 - 2 < 0 \wedge x_2 - 1 < 0$ in gray, and its borders in blue.

Definition 2.4. A *cell* is a set $c \subseteq \mathbb{R}_{\geq 0}^{\mathcal{X}}$, defined by a conjunction of affine inequalities $I_1 \wedge \cdots \wedge I_m$, such that $\nu \in c$ if and only if for all $1 \leq i \leq m$, $\nu \models I_i$. We write $c = \llbracket I_1 \wedge \cdots \wedge I_m \rrbracket$ in this case.

Cells are convex polyhedra, and the intersection of finitely many cells is a cell. From every affine inequality I , we can extract an affine equality $E(I)$, of identical partial derivatives and additive constant. Then, we call *borders* of a cell $c = \llbracket I_1 \wedge \cdots \wedge I_m \rrbracket$ the affine equalities $E(I_1), \dots, E(I_m)$. The closure \bar{c} of a cell c is obtained by replacing every comparison operator $<$ by \leq in its affine inequalities. Note that regions and zones are particular cases of cells, where borders are of the form $x + b = 0$ or $x - y + b = 0$. An example of cell and its borders is given in Figure 4.

We use the notion of cells in order to describe specific cases of *piecewise affine functions* that we will need in this article. First we use them to describe a partition of the space. Let E be an affine equality of equation $a_1x_1 + \cdots + a_nx_n + b = 0$. We say that $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ is partitioned by E into three cells:

- $c_{<}$, defined by $a_1x_1 + \cdots + a_nx_n + b < 0$;
- $c_{>}$, defined by $a_1x_1 + \cdots + a_nx_n + b > 0$, i.e. $-a_1x_1 - \cdots - a_nx_n - b < 0$;
- $c_{=}$, defined by $a_1x_1 + \cdots + a_nx_n + b = 0$, i.e. the conjunction of $a_1x_1 + \cdots + a_nx_n + b \leq 0$ and $-a_1x_1 - \cdots - a_nx_n - b \leq 0$.

Then, given a set $\mathcal{E} = \{E_1, \dots, E_m\}$ of affine equalities, we denote $c_{<}^j$, $c_{>}^j$ and $c_{=}^j$ the three cells obtained from $E_j \in \mathcal{E}$. For every mapping $\phi: \mathcal{E} \rightarrow \{<, >, =\}$, we define c_ϕ as the cell $c_{\phi(E_1)}^1 \cap \cdots \cap c_{\phi(E_m)}^m$. Every valuation of $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ belongs to some c_ϕ , and if $\phi \neq \phi'$ then $c_\phi \cap c_{\phi'} = \emptyset$: hence, the set of mappings $\{<, >, =\}^{\mathcal{E}}$ provides a partition of $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ into 3^m cells. We say that $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ is partitioned by \mathcal{E} into $m' \in \mathbb{N}$ cells if m' of those 3^m cells are non-empty. We denote $\text{Splits}(m, n)$ the greatest m' over any partition of $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ by m affine equalities. Similarly, a cell $c \subseteq \mathbb{R}_{\geq 0}^{\mathcal{X}}$ is partitioned by \mathcal{E} into at most $\text{Splits}(m, n)$ sub-cells that have non-empty intersection with c . In particular, under the bounded clocks assumption we will partition $[0, M]^{\mathcal{X}}$ instead of $\mathbb{R}_{\geq 0}^{\mathcal{X}}$.

Example 2.5. The $\text{Splits}(2, 2) = 9$ cells that partition $[0, 2]^{\mathcal{X}}$ according to $\mathcal{E} = \{2x_1 + x_2 - 2 = 0, x_2 - 1 = 0\}$ are represented in Figure 5.

As a technical tool, we will need to better understand how the number of cells in a partition grows with respect to the number of equalities used to describe this partition. In

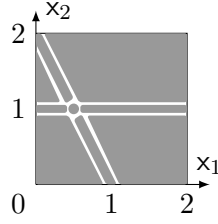


Figure 5: A partition of $[0, 2]^{\mathcal{X}}$ according to the two affine equalities of Figure 4.

fact, the number of cells $\text{Splits}(m, n)$ is bounded by $\mathcal{O}((2m)^n)$, which we will use abundantly to compute the complexity of our algorithms. The technical proof is delayed in Appendix A.

Lemma 2.6. *For all $m, n \geq 0$, $\text{Splits}(m, n) \leq 2^n(m + 1)^n$.*

We summarise our representation of partitions with affine equalities in the following definition:

Definition 2.7. A *partition* P is defined by a cell c_P and a set \mathcal{E}_P of affine equalities, such that P encodes the set of cells that partition c_P according to \mathcal{E}_P , called *base cells*.

The cell c_P is called the domain of P . We denote $[\nu]_P$ the base cell that contains valuation $\nu \in c_P$.

Definition 2.8. A *partition function* F defined over a partition P is a mapping from the base cells of P to affine functions. It encodes a mapping from c_P to \mathbb{R}_∞ , denoted $\llbracket F \rrbracket$: if $\nu \in c_P$ and $F([\nu]_P)$ is defined by $y = a_1x_1 + \dots + a_nx_n + b$, then $\llbracket F \rrbracket(\nu)$ equals $a_1 \cdot \nu(x_1) + \dots + a_n \cdot \nu(x_n) + b$.

A partition function F of domain c_P is *continuous* if for all $\nu \in c_P$, for every base cell c_b such that $\nu \in \bar{c}_b$, if $F(c_b)$ is defined by $y = a_1x_1 + \dots + a_nx_n + b$ then $\llbracket F \rrbracket(\nu) = a_1 \cdot \nu(x_1) + \dots + a_n \cdot \nu(x_n) + b$. In other words, the affine equations provided by F to neighbouring cells should match on the borders that separate them.

Finally, we use a pair (P, F) where P is a partition of domain $[0, M]^{\mathcal{X}}$, and F is a partition function defined over P , to encode a *piecewise affine function* $\llbracket F \rrbracket: [0, M]^{\mathcal{X}} \rightarrow \mathbb{R}_\infty$. The piecewise affine function is said *continuous on regions* if for every region $r \in \text{Reg}(\mathcal{X}, M)$, the restriction of $\llbracket F \rrbracket$ to domain r is continuous. There could be discontinuities in $\llbracket F \rrbracket$, but only at borders separating different regions. In particular, if a partition function is continuous over regions, and $\llbracket F \rrbracket(\nu) = +\infty$ (respectively $-\infty$) for some ν , then for all ν' in the same region as ν , $\llbracket F \rrbracket(\nu') = \llbracket F \rrbracket(\nu)$.

2.4. Weighted timed games. We now turn our attention to (turn-based) weighted timed two-player games with a shortest-path objective towards a set of target locations. We will first define weighted timed games, giving their semantics in terms of an *infinite* transition system on which the traditional notions of strategies and values are defined.

Definition 2.9. A *weighted timed game* (WTG) is a tuple $\mathcal{G} = \langle L_{\text{Min}}, L_{\text{Max}}, \mathcal{X}, L_t, E, \text{wt}, \text{fwt} \rangle$ with $L = L_{\text{Min}} \uplus L_{\text{Max}}$ a finite set of locations split between players Min and Max (in drawings, locations belonging to Min are depicted by circles and the ones belonging to Max by squares), \mathcal{X} a finite set of clocks, $E \subseteq L \times \text{Guards}(\mathcal{X}) \times 2^{\mathcal{X}} \times L$ a finite set of edges $\ell \xrightarrow{g, \mathcal{Y}} \ell'$ from

location ℓ to location ℓ' , labelled by a guard g and a set \mathcal{Y} of clocks to reset, $\mathbf{wt}: E \uplus L \rightarrow \mathbb{Z}$ a weight function associating an integer weight with each location and edge, $L_{\mathbf{t}} \subseteq L_{\text{Min}}$ a set of target locations for player Min, and $\mathbf{fwt}: L_{\mathbf{t}} \times \mathbb{R}_{\geq 0}^{\mathcal{X}} \rightarrow \mathbb{R}_{\infty}$ is a function⁴ mapping each target configuration to a final weight of \mathbb{R}_{∞} .

The semantics of a weighted timed game \mathcal{G} is defined in terms of an infinite labelled transition system $\llbracket \mathcal{G} \rrbracket$ whose states are configurations $(\ell, \nu) \in L \times \mathbb{R}_{\geq 0}^{\mathcal{X}}$. Configurations are split into players according to the location ℓ . A configuration (ℓ, ν) is a target if $\ell \in L_{\mathbf{t}}$, and its final weight is $\mathbf{fwt}(\ell, \nu)$. The labels of $\llbracket \mathcal{G} \rrbracket$ are given by $\mathbb{R}_{\geq 0} \times E$ and will encode the delay that a player wants to spend in the current location, before firing a certain edge. For every delay $d \in \mathbb{R}_{\geq 0}$, edge $e = \ell \xrightarrow{g, \mathcal{Y}} \ell' \in E$ and valuation ν , there is a transition $(\ell, \nu) \xrightarrow{d, e} (\ell', \nu')$ if $\nu + d \models g$ and $\nu' = (\nu + d)[\mathcal{Y} := 0]$. The weight of such a transition takes into account both discrete and continuous costs, and is given by $d \cdot \mathbf{wt}(\ell) + \mathbf{wt}(e)$.

Without loss of generality (since we can detect such deadlocks using classical attractor techniques using regions we describe later, and add a new transition towards a sink location in this case), we suppose the absence of deadlocks in $\llbracket \mathcal{G} \rrbracket$ except on target locations:

Hypothesis 1. For each location $\ell \in L \setminus L_{\mathbf{t}}$ and valuation $\nu \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$, there exist $d \in \mathbb{R}_{\geq 0}$ and $e \in E$ such that $(\ell, \nu) \xrightarrow{d, e} (\ell', \nu')$, and no edges start from $L_{\mathbf{t}}$.

Plays are maximal paths in the transition system $\llbracket \mathcal{G} \rrbracket$: thanks to the previous hypothesis, they are either infinite or end in a target location. First and last elements of a finite play ρ are denoted $\mathbf{first}(\rho)$ and $\mathbf{last}(\rho)$, respectively. For a player $\mathbf{P} \in \{\text{Min}, \text{Max}\}$, the set of non-maximal plays ρ (often called *finite plays*) such that $\mathbf{last}(\rho) \in L_{\mathbf{P}}$ is denoted $\text{FPlays}^{\mathbf{P}}$.

A *strategy* $\sigma_{\mathbf{P}}$ for player \mathbf{P} is a mapping $\text{FPlays}^{\mathbf{P}} \rightarrow \mathbb{R}_{\geq 0} \times E$, such that for all $\rho \in \text{FPlays}^{\mathbf{P}}$ ending in configuration (ℓ, ν) , the transition system $\llbracket \mathcal{G} \rrbracket$ contains a transition labelled by $\sigma_{\mathbf{P}}(\rho)$ from (ℓ, ν) . A strategy is said *positional* (or *memoryless*) if for all $\rho, \rho' \in \text{FPlays}^{\mathbf{P}}$ ending in the same configuration, $\sigma_{\mathbf{P}}(\rho) = \sigma_{\mathbf{P}}(\rho')$. Let $\mathbf{play}((\ell_0, \nu_0), \sigma_{\text{Min}}, \sigma_{\text{Max}})$ denote the unique maximal play starting from configuration (ℓ_0, ν_0) such that for every prefix ρ of $\mathbf{play}((\ell_0, \nu_0), \sigma_{\text{Min}}, \sigma_{\text{Max}})$ in $\text{FPlays}^{\mathbf{P}}$, the next transition in $\mathbf{play}((\ell_0, \nu_0), \sigma_{\text{Min}}, \sigma_{\text{Max}})$ is labelled by $\sigma_{\mathbf{P}}(\rho)$.

The objective of Min is to reach a target location, while minimising the cumulative weight up to the target. Hence, we associate to every finite play $\rho = (\ell_0, \nu_0) \xrightarrow{d_1, e_1} (\ell_1, \nu_1) \cdots \xrightarrow{d_k, e_k} (\ell_k, \nu_k)$ its cumulative weight

$$\mathbf{wt}_{\Sigma}(\rho) = \sum_{i=1}^k d_i \cdot \mathbf{wt}(\ell_{i-1}) + \mathbf{wt}(e_i)$$

Then, the weight of a maximal play ρ , also denoted by $\mathbf{wt}(\rho)$, is defined by $+\infty$ if ρ is infinite and thus does not reach $L_{\mathbf{t}}$, and $\mathbf{wt}_{\Sigma}(\rho) + \mathbf{fwt}(\ell, \nu)$ if it is finite and ends in (ℓ, ν) with $\ell \in L_{\mathbf{t}}$. Then, the respective values of the strategies are defined by

$$\begin{aligned} \mathbf{Val}_{\mathcal{G}}((\ell_0, \nu_0), \sigma_{\text{Min}}) &= \sup_{\sigma_{\text{Max}}} \mathbf{wt}(\mathbf{play}((\ell_0, \nu_0), \sigma_{\text{Min}}, \sigma_{\text{Max}})) \\ \mathbf{Val}_{\mathcal{G}}((\ell_0, \nu_0), \sigma_{\text{Max}}) &= \inf_{\sigma_{\text{Min}}} \mathbf{wt}(\mathbf{play}((\ell_0, \nu_0), \sigma_{\text{Min}}, \sigma_{\text{Max}})) \end{aligned}$$

⁴We restrict the type of functions allowed in Hypothesis 4: informally, we will only deal with piecewise affine functions, as defined previously.

Finally, we let

$$\underline{\text{Val}}(\ell_0, \nu_0) = \sup_{\sigma_{\text{Max}}} \text{Val}_{\mathcal{G}}((\ell_0, \nu_0), \sigma_{\text{Max}}) \quad \overline{\text{Val}}(\ell_0, \nu_0) = \inf_{\sigma_{\text{Min}}} \text{Val}_{\mathcal{G}}((\ell_0, \nu_0), \sigma_{\text{Min}})$$

be the *lower* and *upper values* of configuration (ℓ_0, ν_0) , respectively. We may easily show that $\underline{\text{Val}} \leq \overline{\text{Val}}$. We say that a strategy σ_{Min}^* of Min is ε -*optimal* if, for all initial configurations (ℓ_0, ν_0)

$$\text{Val}_{\mathcal{G}}((\ell_0, \nu_0), \sigma_{\text{Min}}^*) \leq \overline{\text{Val}}(\ell_0, \nu_0) + \varepsilon.$$

It is said *optimal* if this holds for $\varepsilon = 0$. A symmetric definition holds for optimal strategies of Max.

We now impose some more hypotheses on the weighted timed games we consider in this article. First, by mimicking the transformation used for timed automata [BPDG98, Lemma 5], we can turn all diagonal guards into non-diagonal ones by keeping a bit of information for each pair of clocks in the locations. Moreover, by multiplying each rational number in guards by an appropriate positive integer, while doing the same transformation on weights of transitions, we can turn every constant in guards into integers. We therefore assume in the rest of this article that

Hypothesis 2. All guards of WTGs are non-diagonal and contain only integers constants.

Seminal works in weighted timed games [ABM04, BCFL04] have assumed that clocks are *bounded*. This is known to be without loss of generality for (weighted) timed automata [BFH⁺01, Theorem 2]: it suffices to replace transitions with unbounded delays with self-loop transitions periodically resetting the clocks. We do not know if it is the case for the weighted timed games defined above. Indeed, the technique of [BFH⁺01] cannot be directly applied. This would give too much power to player Max that would then be allowed to loop in a location where an unbounded delay could originally be taken before going to the target. In [BCFL04], the situation is simpler since the game is *concurrent*, and thus Min always has a chance to move outside of such a situation. Trying to detect and avoid such situations in our turn-based case seems difficult in the presence of negative weights, since the opportunities of Max crucially depend on the configurations of value $-\infty$ that Min could control afterwards: we will see in Proposition 3.12 that detecting such configurations is undecidable, which is an additional evidence to motivate the decision to focus only on bounded weighted timed games. We thus suppose from now on that

Hypothesis 3. The WTGs are *bounded*, *i.e.* $\llbracket \mathcal{G} \rrbracket$ is restricted to configurations in $L \times [0, M)^{\mathcal{X}}$ where $M \in \mathbb{N}_{>0}$ is the greatest constant appearing in guards.

We need to be able to represent finitely the final weight functions of weighted timed games. The simplest assumption, consisting in somehow disallowing final weights as usually done in the literature, would be to map every target location to the weight 0. We keep more complex final weights, since we will need them in the process of solving weighted timed games. However, it will be sufficient to assume that:

Hypothesis 4. The final weight functions fwt are described by pairs (P, F) where P is a partition of domain $[0, M)^{\mathcal{X}}$, and F is a partition function defined over P : they are thus piecewise affine functions with a finite number of pieces (that are cells, as defined in Definition 2.4), and are assumed to be continuous on regions.

In particular, infinite final weights are constant over regions, *i.e.* if some configuration (ℓ_t, ν) has final weight $+\infty$ or $-\infty$, then for every valuation ν' in the same region as ν ,

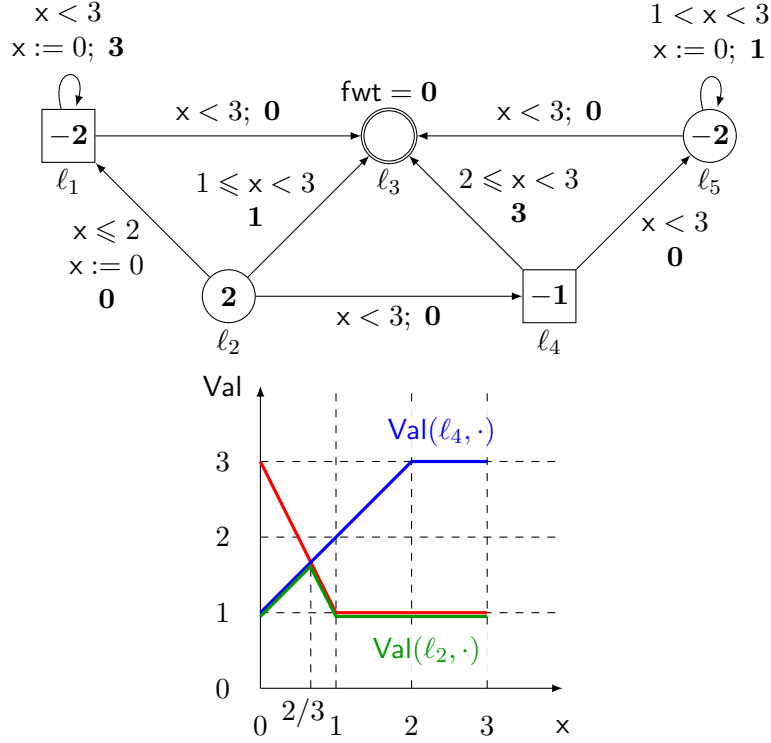


Figure 6: A weighted timed game with a single clock x , and a depiction of its value function. Weights are indicated in bold font on locations and edges. The target location is l_3 , whose final weight function is zero. The blue curve is the values associated to location l_4 , the green curve the values of location l_2 , obtained as the minimum between the blue and the red curve, depicting the weight obtained from l_2 when Min decides to jump directly from l_2 to l_3 .

$fwt(\ell_t, \nu) = fwt(\ell_t, \nu')$. The zero final weight function satisfies this property. Moreover, the computations we will perform in the following maintain this property as an invariant.

Example 2.10. An example of WTG satisfying the previous hypotheses is depicted on Figure 6. Observe that location l_1 (respectively l_5) has value $+\infty$ (respectively $-\infty$). Indeed, from any configuration (ℓ_1, ν) with $\nu \in [0, 3)^x$, player Max can play the self-loop on l_1 , ensuring that the target l_3 is never reached. Moreover, from any configuration (ℓ_5, ν) with $\nu \in [0, 3)^x$, player Min can play the self-loop on l_5 , leading to a reset of clock x . From $(\ell_5, (x = 0))$ Min can pick a delay of 1.5, and loop again while accumulating a weight of $-2(1.5) + 1 = -2$. This loop of l_5 can be iterated arbitrarily many times before playing the edge to l_3 , ensuring a weight arbitrarily low, hence every configuration on l_5 (in the bounded valuation domain $[0, 3)^x$) has value $-\infty$. As a consequence, the value in l_4 is only determined by the transition to l_3 (since Max tries to avoid the target), and depicted in blue in Figure 6. In location l_2 , the value that Min can get while jumping through the transition to l_3 is depicted in red. While jumping to l_2 after a delay 0 (since $wt(\ell_2) > 0$, Min wants to minimise the time spent in this location), Min can also obtain the value of l_4 . Therefore, the value of location l_2 is obtained as the minimum of these two curves, depicted

in green. Observe the intersection point in $x = 2/3$ requiring to refine the usual regions (of granularity 1).

It is known that (turn-based) weighted timed games are determined⁵, *i.e.* $\underline{\text{Val}}(\ell, \nu) = \overline{\text{Val}}(\ell, \nu)$ for each location ℓ and valuation ν , therefore we use the notation $\text{Val}_{\mathcal{G}}$ to refer to both values.

We denote by w_{\max}^L (respectively w_{\max}^E) the maximal weight in absolute values of locations (respectively of edges) in \mathcal{G} :

$$w_{\max}^L = \max_{\ell \in L} |\text{wt}(\ell)| \quad \text{and} \quad w_{\max}^E = \max_{e \in E} |\text{wt}(e)|$$

Moreover, we denote by w_{\max} a bound on the weight of transitions in $\llbracket \mathcal{G} \rrbracket$, that exists since clocks are bounded by M :

$$w_{\max} = Mw_{\max}^L + w_{\max}^E$$

The integer w_{\max} is at most exponential in the size of \mathcal{G} , and can thus be stored in polynomial space.

We consider the *value problem*: given a WTG \mathcal{G} , a location ℓ and a threshold $\alpha \in \mathbb{Q}$, we want to know whether $\text{Val}_{\mathcal{G}}(\ell, \mathbf{0}) \leq \alpha$. We also consider the *$+\infty$ -value problem* (respectively *$-\infty$ -value problem*) asking if $\text{Val}_{\mathcal{G}}(\ell, \mathbf{0})$ equals $+\infty$ (respectively $-\infty$). In the context of timed games, optimal strategies may not exist, even for finite values.⁶ We thus generally focus on the search for ε -optimal strategies, that guarantee the optimal value, up to a small error $\varepsilon \in \mathbb{R}_{>0}$: this is the *synthesis problem*. Moreover, when the value problem is undecidable, we also consider the *value approximation problem* that consists, given a precision $\varepsilon \in \mathbb{Q}_{>0}$, in computing an ε -approximation of $\text{Val}_{\mathcal{G}}(\ell, \mathbf{0})$. More generally, we will try to compute an ε -approximation of the whole value function (and not only for an initial configuration with all clocks being 0): this means that we want to compute (in the format of a pair (P, F) with P a partition of domain $[0, M]^{\mathcal{X}}$ and F a partition function, since we will show that this is sufficient) a function $\mathbf{V}: L \times [0, M]^{\mathcal{X}} \rightarrow \mathbb{R}_{\infty}$ such that $\|\text{Val}_{\mathcal{G}} - \mathbf{V}\|_{\infty} \leq \varepsilon$, where $\|\cdot\|_{\infty}$ denotes the classical ∞ -norm of mappings, so that $\|f\|_{\infty} = \sup_x |f(x)|$. In particular, we ask that $\text{Val}_{\mathcal{G}} = \mathbf{V}$ on all configurations where either function has infinite value, and as such solving the infinite value problems will be a requirement.

For the purpose of stating complexity results we assume that the size needed to encode an input WTG $\mathcal{G} = \langle L_{\text{Min}}, L_{\text{Max}}, \mathcal{X}, L_{\text{t}}, E, \text{wt}, \text{fwt} \rangle$ is linear in $|L|$, $|E|$, and $n = |\mathcal{X}|$. We assume that integer constants are encoded in binary, so that the input depends logarithmically in w_{\max}^L , w_{\max}^E and M . Finally, we assume that the final weight function fwt is encoded as a partition function (P, F) , so that the input is linear in the number of affine expressions used in P and F and logarithmic in the size of their constants. We encode all rational constants as irreducible fractions of binary integers.

2.5. Related work. In the one-player case, computing the optimal value and an ε -optimal strategy for WTGs (called weighted or priced timed automata) is known to be PSPACE-complete [BBBR07]. In the two-player case, the value problem of WTGs (also called priced timed games in the literature) is undecidable with 3 clocks [BBR05, BJM15], or even 2

⁵The result is stated in [BGH⁺15] for weighted timed games (called priced timed games) with one clock, but the proof does not use the assumption on the number of clocks.

⁶For example, a player may want to let time elapse as much as possible, but with delay $d < 1$ because of a strict guard.

clocks in the presence of negative weights [BGNK⁺14] (for the existence problem, asking if a strategy of player Min can guarantee a given threshold).

To obtain decidability, one possibility has been to limit the number of clocks to 1: then, there is an exponential-time algorithm to compute the value as well as ε -optimal strategies in the presence of non-negative weights only [BBM06, Rut11, HIJM13], whereas the problem is only known to be PTIME-hard. A similar result can be lifted to arbitrary weights, under restrictions on the resets of the clock in cycles [BGH⁺15].

The other possibility to obtain a decidability result [BCFL04, ABM04] with non-negative weights only is to enforce a semantical property of divergence, originally called strictly non-Zeno cost: it asks that every play following a cycle of the region abstraction (see later for a formal definition) has weight at least 1. In [BJM15], the authors slightly extend the strictly non-Zeno cost property, to allow for cycles of weight exactly 0 while still preventing those of weight arbitrarily close to 0. Unfortunately, this leads to an undecidable value problem, but they propose a solution to the value approximation problem. **This article aims at extending these two restrictions in the presence of both non-negative and negative weights calling it the divergence property, in order to obtain either decidability of the value problem or approximations of the value for a large class of multi-clocks weighted timed games in the presence of arbitrary weights.**

Other objectives, not directly related to optimal reachability, have been considered in [BCR14] for WTG, like mean-payoff and parity objectives. In this work, the authors manage to solve these problems for the so-called class of δ -robust WTGs that they introduce.⁷

2.6. Region abstraction. The partition of the configuration space into regions can be maintained throughout the play in a WTG. By forgetting about the precise valuation of clocks, we obtain a *region abstraction* (that is usually called the *region automaton* in the literature).

Definition 2.11. Given a WTG $\mathcal{G} = \langle L_{\text{Min}}, L_{\text{Max}}, \mathcal{X}, L_{\text{t}}, E, \text{wt}, \text{fwt} \rangle$ such that all clocks are bounded by M and all guards belong to $\text{Guards}_N^{\text{nd}}(\mathcal{X}, M)$ for some granularity $1/N$, we define the *region abstraction* of \mathcal{G} as the finite labelled transition system $\langle L \times \text{Reg}_N(\mathcal{X}, M), T \rangle$ labelled over $\text{Reg}_N(\mathcal{X}, M) \times E$, where T contains all transitions $(\ell, r) \xrightarrow{r'', e} (\ell', r')$ such that $e = \ell \xrightarrow{g, \mathcal{Y}} \ell'$ is an edge of \mathcal{G} , r'' is a time-successor of r , $r'' \models g$ and $r''[\mathcal{Y} := 0] = r'$.

The states of the region abstraction are called *region states*, and its paths are called *region paths*. As there are finitely many regions, the region abstraction of \mathcal{G} is a finite transition system, where paths \mathbf{p} represent sequences of regions alternating between letting time elapse and taking edges. From a play $\rho = (\ell_0, \nu_0) \xrightarrow{d_1, e_1} (\ell_1, \nu_1) \xrightarrow{d_2, e_2} \dots$ in \mathcal{G} , we construct a region path $\mathbf{p} = (\ell_0, [\nu_0]) \xrightarrow{[\nu_0 + d_1], e_1} (\ell_1, [\nu_1]) \xrightarrow{[\nu_1 + d_2], e_2} \dots$, and say that ρ follows \mathbf{p} .

From the region abstraction, we can construct a *region game* that can be seen as a product of the original WTG with the region abstraction.

Definition 2.12. Given a WTG $\mathcal{G} = \langle L_{\text{Min}}, L_{\text{Max}}, \mathcal{X}, L_{\text{t}}, E, \text{wt}, \text{fwt} \rangle$ such that all clocks are bounded by M and all guards belong to $\text{Guards}_N^{\text{nd}}(\mathcal{X}, M)$ for some granularity $1/N$, we define the *region game* of \mathcal{G} as the WTG $\mathcal{R}_N(\mathcal{G}) = \langle L_{\text{Min}} \times \text{Reg}_N(\mathcal{X}, M), L_{\text{Max}} \times \text{Reg}_N(\mathcal{X}, M), \mathcal{X}, L_{\text{t}} \times \text{Reg}_N(\mathcal{X}, M), E', \text{wt}', \text{fwt}' \rangle$ whose locations are region states, wt' and fwt' are obtained trivially

⁷As mentioned in introduction, while our divergent games have a similar definition, the two classes are incomparable.

from wt and fwt by forgetting about regions, and E' is defined by transforming every transition $(\ell, r) \xrightarrow{r'', e} (\ell', r')$ in T , where e is labelled by (g, \mathcal{Y}) , into an edge $(\ell, r) \xrightarrow{g'', \mathcal{Y}} (\ell', r')$, with g'' a guard chosen arbitrarily so that $\llbracket g'' \rrbracket = r'' \subseteq \llbracket g \rrbracket$.

Every play in \mathcal{G} exists in $\mathcal{R}_N(\mathcal{G})$ as a play following a region path \mathbf{p} , and conversely every play in $\mathcal{R}_N(\mathcal{G})$ is a valid play in \mathcal{G} , by projecting away the region information of $\mathcal{R}_N(\mathcal{G})$. We thus obtain:

Lemma 2.13. *Games \mathcal{G} and $\mathcal{R}_N(\mathcal{G})$ contain the same plays. For all $\ell \in L$, $1/N$ -regions r , and $\nu \in r$, $\text{Val}_{\mathcal{G}}(\ell, \nu) = \text{Val}_{\mathcal{R}_N(\mathcal{G})}((\ell, r), \nu)$.*

As we assume that guards have integer constants, we can use the granularity $1/N = 1$, and we will write $\mathcal{R}(\mathcal{G})$ instead of $\mathcal{R}_N(\mathcal{G})$ in that case. Finally, we denote by $|\mathcal{R}(\mathcal{G})|$ the number of locations in the region game, equal to $|L| |\text{Reg}(\mathcal{X}, M)|$.

2.7. Corner-point abstraction. Despite all the interest and success of regions to study timed systems, they are not sufficient to handle weighted timed automata/games. Indeed, for a single-clock case, and in a location ℓ of weight 1, spending time in ℓ from region $(0, 1)$ to region $\{2\}$ can cost any possible weight in the interval $(1, 2)$. We therefore must also keep a more precise information of *where we are inside each region*. This is the goal of the *corner-point abstraction* introduced in [BFH⁺01, LBB⁺01] to study one-player WTGs with non-negative weights, generalised in [BBBR07] to handle negative weights, and in [BBL08] for the multi-cost setting.

If r is an $1/N$ -region, let \bar{r} denote its topological closure, *i.e.* the smallest zone that contains r associated to a non-strict guard. The *corners* of r are all the valuations in \bar{r} that belong to $\mathbb{Q}_{\geq 0}^{\mathcal{X}}$. If r is characterised by (ι, β) with β the partition $\beta_0 \uplus \beta_1 \uplus \dots \uplus \beta_m$, then ι is a corner of r . If $m = 0$, then $r = \{\iota\}$, otherwise r does not include its corners but contains valuations arbitrarily close to them. The corners of r are the vertices of the polytope \bar{r} , such that \bar{r} is their convex hull. There are at most $n + 1$ corners in each $1/N$ -region. The corners of the green region in Figure 3 are the valuations $(1, 0)$ and $(1, 1)$.

We call corner state a triple (ℓ, r, \mathbf{v}) that contains information about a region state (ℓ, r) of $\mathcal{R}_N(\mathcal{G})$, and a corner \mathbf{v} of the $1/N$ -region r .

Notice that reset operations preserve the corners, *i.e.*, if \mathbf{v} is a corner of the region r , then $\mathbf{v}[x := 0]$ is a corner of the region $r[x := 0]$. This allows one to enrich the region game with corner information:

Definition 2.14. The corner-point abstraction $\Gamma_N(\mathcal{G})$ of a WTG \mathcal{G} is the WTG obtained as a refinement of $\mathcal{R}_N(\mathcal{G})$ where guards on edges are enforced to stay on one of the corners of the current $1/N$ -region: the locations of $\Gamma_N(\mathcal{G})$ are all corner states of $\mathcal{R}_N(\mathcal{G})$, associated to each player accordingly, and edges are all $(\ell, r, \mathbf{v}) \xrightarrow{g'', \mathcal{Y}} (\ell', r', \mathbf{v}')$ such that there exists a region r'' , an edge $t = (\ell, r) \xrightarrow{g, \mathcal{Y}} (\ell', r')$ of $\mathcal{R}_N(\mathcal{G})$ such that the model of guard g'' is a corner \mathbf{v}'' of region r'' satisfying the guard \bar{g} (recall that \bar{g} is the closed version of g), $\mathbf{v}'' \in \text{PostTime}(\mathbf{v})$, $\mathbf{v}' = \mathbf{v}''[\mathcal{Y} := 0]$, $r' = r''[\mathcal{Y} := 0]$ (the two last conditions ensure that \mathbf{v}' is indeed a corner of r') and there exist two valuations $\nu \in r$, $\nu' \in r'$ such that $((\ell, r), \nu) \xrightarrow{d', t} ((\ell', r'), \nu')$ for some $d' \in \mathbb{R}_{\geq 0}$ (the latter condition ensures that the edge between corners is not spurious, *i.e.* created by the closure of guards). Weights of locations and edges are trivially recovered from $\Gamma_N(\mathcal{G})$. We define the final weight function of $\Gamma_N(\mathcal{G})$ over the only valuation \mathbf{v} reachable

in location (ℓ, r, \mathbf{v}) (with $\ell \in L_t$) by $\text{fwt}((\ell, r, \mathbf{v}), \mathbf{v}) = \lim_{\nu \rightarrow \mathbf{v}, \nu \in r} \text{fwt}(\ell, \nu)$ (the limit is well defined since fwt is piecewise affine with a finite number of pieces on region r , by Hypothesis 4).

We denote by $|\Gamma_N(\mathcal{G})|$ the number of locations in the corner-point abstraction, bounded by $|L| |\text{Reg}_N(\mathcal{X}, M)| (n + 1)$.

The WTG $\Gamma_N(\mathcal{G})$ can be seen as a finite weighted game, *i.e.* a WTG without clocks, by removing guards, resets and rates of locations, and replacing the weights of edges by the actual weight of jumping from one corner to another: an edge $((\ell, r), \mathbf{v}) \xrightarrow{g'', \mathcal{Y}} ((\ell', r'), \mathbf{v}')$ becomes a transition from $((\ell, r), \mathbf{v})$ to $((\ell', r'), \mathbf{v}')$ with weight $d \cdot \text{wt}(\ell) + \text{wt}(t)$, with $d \in \mathbb{R}_{\geq 0}$ the only delay such that $\llbracket g'' \rrbracket = \{\mathbf{v} + d\}$.⁸ Note that delay d is necessarily a rational of the form α/N with $\alpha \in \mathbb{N}$, since it must relate corners of $1/N$ -regions. In particular, this proves that the cumulative weight $\text{wt}_\Sigma(\bar{\rho})$ of a finite play $\bar{\rho}$ in $\Gamma_N(\mathcal{G})$ is indeed a rational number with denominator N .

We will call *corner play* every play $\bar{\rho}$ in the corner-point abstraction $\Gamma_N(\mathcal{G})$: it can also be interpreted as an execution in \mathcal{G} where all guards are closed (as explained in the definition above). It straightforwardly projects on a finite path \mathbf{p} in the region game $\mathcal{R}_N(\mathcal{G})$: in this case, we say again that $\bar{\rho}$ follows \mathbf{p} . Figure 7 depicts a play, its projected path in the region game and one of its associated corner plays.

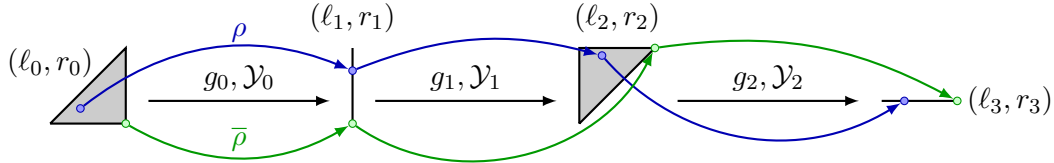


Figure 7: A play ρ (in blue), its projected path \mathbf{p} in the region game (in black), and one of its associated corner plays $\bar{\rho}$ (in green).

Let $\bar{\rho}$ be a corner play following a region path \mathbf{p} . The weight of $\bar{\rho}$ refers to its weight in $\Gamma_N(\mathcal{G})$. It is possible to find a play ρ following \mathbf{p} close to $\bar{\rho}$, in the sense that we control the difference between their respective cumulative weights:

Lemma 2.15 [BBL08, Prop. 5]. *For all $\varepsilon > 0$, all finite region paths \mathbf{p} , and all corner plays $\bar{\rho}$ following \mathbf{p} , there exists a play ρ in \mathcal{G} following \mathbf{p} such that $|\text{wt}_\Sigma(\rho) - \text{wt}_\Sigma(\bar{\rho})| \leq \varepsilon$.*

Thus, corner plays allow one to obtain faithful information on the plays that follow the same path.

Lemma 2.16. *If \mathbf{p} is a finite region path in $\mathcal{R}_N(\mathcal{G})$, the set of cumulative weights $\{\text{wt}_\Sigma(\rho) \mid \rho \text{ play of } \mathcal{G} \text{ following } \mathbf{p}\}$ is an interval bounded by the minimum and the maximum values of the set $\{\text{wt}_\Sigma(\bar{\rho}) \mid \bar{\rho} \text{ corner play of } \Gamma_N(\mathcal{G}) \text{ following } \mathbf{p}\}$.*

Proof. The set $\{\text{wt}_\Sigma(\rho) \mid \rho \text{ finite play following } \mathbf{p}\}$ is an interval as the image of a convex set by an affine function (see [BBBR07, Sec. 3.2] for an explanation).

The good properties of the corner-point abstraction allow us to conclude, since for every play ρ following \mathbf{p} , one can find a corner play following \mathbf{p} of smaller weight and one of larger

⁸In case several edges lead to the same transition, for instance when two transitions with different guards reset all clocks, we either allow for multi-transitions or choose the best weight according to the player owning the current location.

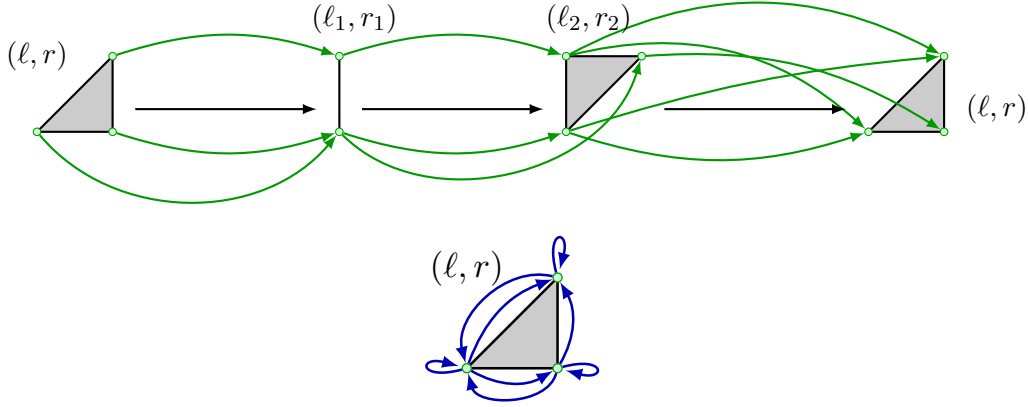


Figure 8: A region cycle \mathbf{p} in the region game (in black), its associated corner plays (in green), and its folded orbit graph (in blue). Note that there is no edge between the top right corner and the bottom right corner, as no corner play goes from the former to the latter.

weight [BFH⁺01, Lemma 1], and for every corner play ρ following \mathbf{p} and every $\varepsilon > 0$, one can find a play following \mathbf{p} whose weight is at most ε away from $\text{wt}_\Sigma(\rho)$ by Lemma 2.15. \square

An important property of the corner-point abstraction, derived from the assumption on the absence of deadlocks in the game, is that corner plays cannot get stuck as long as they follow a region path:

Lemma 2.17 [Pur00, Lem. 8]. *Let \mathbf{p} be a region path starting from (ℓ, r) and ending in (ℓ', r') . For all corners \mathbf{v} of r , there exists a corner play following \mathbf{p} that starts in (ℓ, r, \mathbf{v}) . For all corners \mathbf{v}' of r' there exists a corner play following \mathbf{p} that ends in (ℓ', r', \mathbf{v}') .*

Useful theoretical tools stem from the corner-point abstraction. Notably, let us focus on a cycle of the region abstraction. In order to study some properties of the corner plays following this cycle, we only need to consider the aggregation of all the behaviours following it. Inspired by the *folded orbit graphs* (FOG) introduced in [Pur00], we define the folded orbit graph $\text{FOG}(\mathbf{p})$ of a region cycle $\mathbf{p} = (\ell_1, r = r_1) \xrightarrow{e_1} (\ell_2, r_2) \xrightarrow{e_2} \dots \xrightarrow{e_k} (\ell_1, r)$ in $\mathcal{R}_N(\mathcal{G})$ as a graph whose vertices are the corners states of region r , and that contains an edge from corner \mathbf{v} to corner \mathbf{v}' if there exists a corner play $\bar{\rho}$ from (ℓ_1, r, \mathbf{v}) to (ℓ_1, r, \mathbf{v}') following \mathbf{p} . We fix $\bar{\rho}$ arbitrarily and label the edge between \mathbf{v} and \mathbf{v}' in $\text{FOG}(\mathbf{p})$ by this corner play: it is then denoted by $\mathbf{v} \xrightarrow{\bar{\rho}} \mathbf{v}'$. An example is depicted in Figure 8.

The folded orbit graph inherits interesting topological properties from the corner-point abstraction. Notably, by Lemma 2.17, for all vertices \mathbf{v} , there exists at least one outgoing edge $\mathbf{v} \xrightarrow{\bar{\rho}'} \mathbf{v}'$, and at least one incoming edge $\mathbf{v}'' \xrightarrow{\bar{\rho}''} \mathbf{v}$ in $\text{FOG}(\mathbf{p})$.

2.8. Value iteration algorithm. The value of a game has been defined as a mapping of each configuration (ℓ, ν) to a value in \mathbb{R}_∞ . We call *value functions* such mappings from $L \times \mathbb{R}_{\geq 0}^X$ to \mathbb{R}_∞ . If \mathbf{V} represents a value function, we denote by \mathbf{V}_ℓ the mapping $\nu \mapsto \mathbf{V}(\ell, \nu)$. As observed in [BCFL04, ABM04], one step of the game is summarised in the following operator \mathcal{F} mapping each value function \mathbf{V} to a value function $\mathbf{V}' = \mathcal{F}(\mathbf{V})$ defined by

$\mathbf{V}'_\ell(\nu) = \text{fwt}(\ell, \nu)$ if $\ell \in L_t$, and otherwise

$$\mathbf{V}'_\ell(\nu) = \begin{cases} \sup_{(\ell, \nu) \xrightarrow{d,e} (\ell', \nu')} [d \cdot \text{wt}(\ell) + \text{wt}(e) + \mathbf{V}_{\ell'}(\nu')] & \text{if } \ell \in L_{\text{Max}} \\ \inf_{(\ell, \nu) \xrightarrow{d,e} (\ell', \nu')} [d \cdot \text{wt}(\ell) + \text{wt}(e) + \mathbf{V}_{\ell'}(\nu')] & \text{if } \ell \in L_{\text{Min}} \end{cases} \quad (2.1)$$

where $(\ell, \nu) \xrightarrow{d,e} (\ell', \nu')$ ranges over valid transitions in \mathcal{G} .

Then, starting from \mathbf{V}^0 mapping every configuration (ℓ, ν) to $+\infty$, except for the targets mapped to $\text{fwt}(\ell, \nu)$, we let $\mathbf{V}^i = \mathcal{F}(\mathbf{V}^{i-1})$ for all $i > 0$. The value function \mathbf{V}^i contains the value $\text{Val}_{\mathcal{G}}^i$, which is intuitively what Min can guarantee when forced to reach the target in at most i steps. More formally, we define $\text{wt}^i(\rho)$ the weight of a maximal play ρ at horizon i , as $\text{wt}(\rho)$ if ρ reaches a target in at most i steps, and $+\infty$ otherwise. Then, $\text{Val}_{\mathcal{G}}^i(\ell, \nu) = \inf_{\sigma_{\text{Min}}} \sup_{\sigma_{\text{Max}}} \text{wt}^i(\text{play}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}}))$ refers to the value at horizon i .

Remark 2.18. In case of non-weighted games, with or without time, the value iteration algorithm is generally called *attractor*. Starting from $+\infty$ for every configuration, except for the target mapped to 0, we compute the previous iterates (with all weights mapped to 0). In this case, as shown in [AMPS98], values 0 or $+\infty$ stay constant over each regions, and there are thus a finite number of possible functions, which ensures that the computation ends in finite time.

We compare value functions componentwise: if \mathbf{V}, \mathbf{V}' are two value functions, we let $\mathbf{V} \leq \mathbf{V}'$ if $\mathbf{V}(\ell, \nu) \leq \mathbf{V}'(\ell, \nu)$ for all configurations (ℓ, ν) . Notice that \mathcal{F} is a monotonic operator, *i.e.* if $\mathbf{V} \leq \mathbf{V}'$, then $\mathcal{F}(\mathbf{V}) \leq \mathcal{F}(\mathbf{V}')$. Moreover, $\mathcal{F}(\mathbf{V}^0) \leq \mathbf{V}^0$ since \mathbf{V}^0 maps every non-target state to $+\infty$, and target states keep the same value. It follows that the sequence $(\mathbf{V}^i)_{i \in \mathbb{N}}$ is non-increasing, as $\mathbf{V}^i = \mathcal{F}^i(\mathbf{V}^0) \geq \mathcal{F}^i(\mathcal{F}(\mathbf{V}^0)) = \mathbf{V}^{i+1}$.

Let us now present known results for the special case of games with no clocks. In this case, the definition of the operator \mathcal{F} of (2.1) is simplified into

$$\mathbf{V}'_\ell = \begin{cases} \min_{\ell \rightarrow \ell'} [\text{wt}(\ell, \ell') + \mathbf{V}_{\ell'}] & \text{if } \ell \in L_{\text{Min}} \\ \max_{\ell \rightarrow \ell'} [\text{wt}(\ell, \ell') + \mathbf{V}_{\ell'}] & \text{if } \ell \in L_{\text{Max}} \end{cases} \quad (2.2)$$

Notice that we remove the valuation part of the notation: configurations are thus simply locations of the game. The value iteration algorithm proposed in [BGHM16] consists in finding the greatest fixpoint of operator \mathcal{F} , *i.e.* the limit of the sequence $(\mathbf{V}^i)_{i \in \mathbb{N}}$. Indeed, this greatest fixpoint is known to be the vector of values of the game (see, *e.g.*, [BGHM16, Corollary 11]). For the special case of acyclic games of depth d , the fixpoint is reached after d steps, and $\text{Val} = \mathbf{V}^d$. In this case, the infinite values in \mathcal{G} (*i.e.* configurations ℓ with $\text{Val}_{\mathcal{G}}(\ell) \in \{-\infty, +\infty\}$) are derived from reaching targets with infinite final weights. If the game contains cycles, infinite values can also come from arbitrarily long plays: a state can have value $+\infty$ if Max can force an infinite play, never reaching any target, and it can have value $-\infty$ if Min can enforce an arbitrarily low weight, *e.g.* by staying in a cycle of negative cumulative weight. These $+\infty$ states correspond to a safety objective for player Max, and can be computed in polynomial time: it is shown in [BGHM16] that for all locations ℓ , $\text{Val}_{\mathcal{G}}(\ell) = +\infty$ if and only if $\mathbf{V}_\ell^{|L|} = +\infty$. In contrast, deciding if a location has value $-\infty$ has no known polynomial solution (it is as hard as solving mean-payoff games). In [BGHM16], it is shown that in the presence of negative weights the sequence $(\mathbf{V}^i)_{i \in \mathbb{N}}$ stabilises after a number of iterations pseudo-polynomial on states with value in $\mathbb{R} \cup \{+\infty\}$, and that states

with value $-\infty$ can be detected in this computation (they are those where the computed value goes under a given threshold).

Such a computation of the greatest fixed point might not be possible in the timed setting, since the value problem is undecidable in general. However, in [ABM04], it is shown that starting from a value function \mathbf{V} that is represented, for each location, by a pair (P, F) where P is partition and F is a partition function defined over P , the same is true for $\mathcal{F}(\mathbf{V})$. Indeed, cells are bounded, convex polyhedra, over which elementary operations (emptiness, intersection and inclusion tests) can be performed with linear programming, and thus in polynomial time. As we will formally recall in Section 7, this data structure allows one to effectively compute the iterates $(\mathbf{V}^i)_{i \in \mathbb{N}}$. Contrary to the case without clocks, recalled before, this sequence does not stabilise in general. The decidability results obtained before (for tree-shaped weighted timed games [ABM04], or strictly non-Zeno weighted timed games [BCFL04], *e.g.*), as well as the ones we obtain in this article, are all based on reasons to either make the sequence stabilise or stop its computation after a sufficient number i of turns to obtain a good approximation of the value.

Remark 2.19. In [ABM04], the domain of partitions is always a single region, and one value function is associated to each region. We define value functions over $[0, M]^{\mathcal{X}}$ instead, in order to obtain a symbolic algorithm, independent of regions. This induces slight differences in the way value functions are defined, because the mappings of [ABM04] are continuous everywhere while ours can have discontinuities at borders between regions. They define their partitions with overlaps over borders, such that $\mathbb{R}_{\geq 0}^{\mathcal{X}}$ is partitioned by an affine equality into two cells, c_{\leq} and c_{\geq} , instead of the three $c_{<}$, $c_{>}$ and $c_{=}$. This changes the number of cells $\text{Splits}(m, n)$ to $\mathcal{O}(m^n)$ instead of $\mathcal{O}((2m)^n)$.

3. DIVERGENT AND ALMOST-DIVERGENT WTGS

In this section, we introduce several classes of weighted timed games for which we state the results that this article shows. All classes are defined in terms of the underlying timed automaton, without making use of the partition of locations into players. Let us start with the class of weighted timed games studied in [BCFL04], to our knowledge the greatest class of WTG where the value problem is known to be decidable.

Definition 3.1. A weighted timed game \mathcal{G} with non-negative weights satisfies the *strictly non-Zeno cost* property when every finite play ρ in \mathcal{G} following a cycle in the region automaton $\mathcal{R}(\mathcal{G})$ satisfies $\text{wt}_{\Sigma}(\rho) \geq 1$.

The intuition behind this class is that the weight of any long enough execution in \mathcal{G} will ultimately grow above any fixed bound, and diverge towards $+\infty$ for an infinite execution. Therefore, the value of \mathcal{G} is equal to the value $\text{Val}_{\mathcal{G}}^i$ at some horizon i large enough (as defined in Section 2.8), making the value problem decidable. It is shown in [BCFL04] that i can be bounded exponentially in the size of \mathcal{G} .

We introduce divergent weighted timed games, as a natural generalisation of the strictly non-Zeno cost property to weights in \mathbb{Z} .

Definition 3.2. A WTG \mathcal{G} is *divergent* when every finite play ρ in \mathcal{G} following a cycle in the region automaton $\mathcal{R}(\mathcal{G})$ satisfies $\text{wt}_{\Sigma}(\rho) \notin (-1, 1)$.

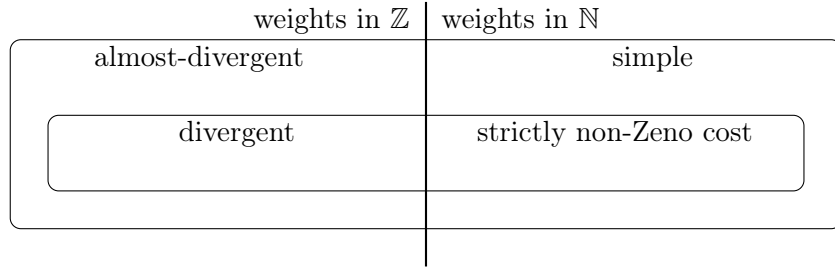


Figure 9: Classes of weighted timed games, and their respective restrictions to non-negative weights.

If \mathcal{G} has only non-negative weights on locations and edges, this definition matches with the strictly non-Zeno cost property of [BCFL04], we will therefore refer to their class as the class of divergent WTG with non-negative weights.

Remark 3.3. As in [BCFL04], we could replace $(-1, 1)$ by $(-\kappa, \kappa)$ to define a notion of κ -divergence. However, since weights and guard constraints in weighted timed games are integers, for $\kappa \in (0, 1)$, a weighted timed game \mathcal{G} is κ -divergent if and only if it is divergent. This will be formally implied by Proposition 4.3 and Lemma 2.16.

Our contributions on divergent WTGs summarise as follows:

Theorem 3.4. *The value problem over divergent WTGs is decidable in 3-EXPTIME, and is EXPTIME-hard. Moreover, deciding if a given WTG is divergent is a PSPACE-complete problem.*

In [BJM15], the authors slightly extend the strictly non-Zeno cost property, to allow for cycles of weight exactly 0 while still preventing those of weight arbitrarily close to 0:

Definition 3.5. A WTG \mathcal{G} with non-negative weights is called *simple* when every finite play ρ in \mathcal{G} following a cycle in the region automaton $\mathcal{R}(\mathcal{G})$ satisfies $\text{wt}_\Sigma(\rho) \in \{0\} \cup [1, +\infty)$.

Unfortunately, it is shown in [BJM15] that the value problem is undecidable for simple WTGs. They propose a solution to the value approximation problem, as a procedure computing an approximation of the value of every configuration. The intuition is that cycles of weight exactly 0 are only possible when every (non-negative) weight encountered along the cycle equals 0, allowing one to define a subgame where every cyclic execution has weight 0. One can then analyse this subgame separately, by applying a semi-unfolding procedure on $\mathcal{R}(\mathcal{G})$.

We introduce a class of WTGs that will extend the notion of simple WTGs and allow negative weights, so that cyclic executions of weight exactly 0 are allowed, but not those close to 0. The first attempt would lead to the requirement that every finite play following a cycle in the region automaton $\mathcal{R}(\mathcal{G})$ has a weight in $(-\infty, -1] \cup \{0\} \cup [1, +\infty)$. However, we did not obtain positive results for the value approximation problem on this class of WTGs since cycles of weight exactly 0 do not have the good property presented above for simple WTGs. Instead, we require a stability by decomposition for cycles of weight 0.

If $\mathbf{p} = (\ell_0, r_0) \xrightarrow{r_1, e_0} (\ell_1, r_1) \xrightarrow{r_2, e_1} \dots (\ell_{k-1}, r_{k-1}) \xrightarrow{r_0, e_{k-1}} (\ell_0, r_0)$ is a region cycle in $\mathcal{R}(\mathcal{G})$, it is either simple (*i.e.* for all i, j such that $0 \leq i < j < k$, $(\ell_i, r_i) \neq (\ell_j, r_j)$) or we can extract smaller cycles from it. Indeed, if \mathbf{p} is not simple, there exists a pair (i, j) such that

$0 \leq i < j < k$ and $(\ell_i, r_i) = (\ell_j, r_j)$. Then, for such a pair, we can write $\mathbf{p} = \mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3$ such that $|\mathbf{p}_1| = i$, $|\mathbf{p}_3| = k - j$. It follows that \mathbf{p}_2 and $\mathbf{p}_3 \mathbf{p}_1$ are both region cycles around (ℓ_i, r_i) . This process is called a decomposition of \mathbf{p} into smaller cycles $\mathbf{p}' = \mathbf{p}_3 \mathbf{p}_1$ and $\mathbf{p}'' = \mathbf{p}_2$.

Definition 3.6. A WTG \mathcal{G} is *almost-divergent* if every play ρ following a cycle \mathbf{p} of $\mathcal{R}(\mathcal{G})$ satisfies $\text{wt}_\Sigma(\rho) \in (-\infty, -1] \cup \{0\} \cup [1, +\infty)$, and if $\text{wt}_\Sigma(\rho) = 0$ then for every decomposition of \mathbf{p} into smaller cycles \mathbf{p}' and \mathbf{p}'' , and plays ρ' and ρ'' following \mathbf{p}' and \mathbf{p}'' , respectively, it holds that $\text{wt}_\Sigma(\rho') = \text{wt}_\Sigma(\rho'') = 0$.⁹

Clearly, every divergent WTG is almost-divergent. Moreover, as we will see in Proposition 4.4, when weights are non-negative, this class matches the simple WTGs of [BJM15] therefore inheriting their undecidability result. We will thus refer to simple WTGs as almost-divergent WTGs with non-negative weights. Figure 9 represents the hierarchy of the classes of WTG that we introduced.

Example 3.7. Consider the WTG \mathcal{G} in Figure 6. The self-loop on ℓ_1 contains a cycle of $\mathcal{R}(\mathcal{G})$ around $(\ell_1, \mathbf{0})$ that jumps to the region $1 < x < 2$. For every $d \in (1, 2)$, there exists a play ρ following this cycle that uses delay d , so that $\text{wt}_\Sigma(\rho) = 3 - 2d \in (-1, 1)$. It follows that \mathcal{G} is neither divergent nor almost-divergent. Changing the guard on this self-loop to $2 \leq x < 3$ makes \mathcal{G} divergent, as every region cycle left in \mathcal{G} iterates the self-loops around $(\ell_1, \mathbf{0})$ and $(\ell_5, \mathbf{0})$ of cumulative weights in $(-3, -1]$ and $(-5, -1)$, respectively.

Example 3.8. Consider the WTG \mathcal{G} in Figure 10, and its region game $\mathcal{R}(\mathcal{G})$. We chose an example where $\mathcal{R}(\mathcal{G})$ is isomorphic to \mathcal{G} for readability reasons. $\mathcal{R}(\mathcal{G})$ contains one SCC $\{\ell_1, \ell_2, \ell_3, \ell_4\}$, made of two simple cycles, $\mathbf{p}_1 = \ell_1 \rightarrow \ell_2 \rightarrow \ell_1$ and $\mathbf{p}_2 = \ell_1 \rightarrow \ell_3 \rightarrow \ell_4 \rightarrow \ell_1$, so that:

- all plays following \mathbf{p}_1 have cumulative weight in the interval $(1, 3)$,
- and all plays following \mathbf{p}_2 have cumulative weight 0. This can be checked by Lemma 2.16.

As every cycle \mathbf{p} of $\mathcal{R}(\mathcal{G})$ either iterates \mathbf{p}_2 only, or contains \mathbf{p}_1 , it holds that a play ρ following \mathbf{p} satisfies $\text{wt}_\Sigma(\rho) \in \{0\} \cup [1, +\infty)$, and if $\text{wt}_\Sigma(\rho) = 0$ then any decomposition of \mathbf{p} into smaller cycles \mathbf{p}' and \mathbf{p}'' implies that they all follow iterates of \mathbf{p}_2 , so that plays along them must have cumulative weight 0. Therefore, \mathcal{G} is almost-divergent. If one removes ℓ_4 , \mathcal{G} becomes divergent.

Our first result on almost-divergent WTGs is the following extension of the approximation procedure for non-negative weights:

Theorem 3.9. *Given an almost-divergent WTG \mathcal{G} , a location ℓ and $\varepsilon \in \mathbb{Q}_{>0}$, we can compute an ε -approximation of $\text{Val}_{\mathcal{G}}(\ell, \mathbf{0})$ in time triply-exponential in the size of \mathcal{G} and polynomial in $1/\varepsilon$. Moreover, deciding if a WTG is almost-divergent is PSPACE-complete.*

To obtain these results on divergent and almost-divergent WTGs, we follow a computation schema that we now outline. First, we will always reason on the region game $\mathcal{R}(\mathcal{G})$ of the almost-divergent WTG \mathcal{G} . The goal is to compute an ε -approximation of $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell_0, [\mathbf{0}]), \mathbf{0})$ for some initial location ℓ_0 . Techniques of [ABM04] (that we will recall in Section 7) allow one to compute the (exact) values of a WTG played on a finite tree, using operator \mathcal{F} of Section 2.8. The idea is thus to decompose as much as possible the game $\mathcal{R}(\mathcal{G})$ as a WTG over a tree. First, we decompose the region game into strongly connected components (SCCs),

⁹Once again, we could replace $-1, 1$ by $-\kappa, \kappa$ with $0 < \kappa < 1$ to define an equivalent notion of κ -almost-divergence.

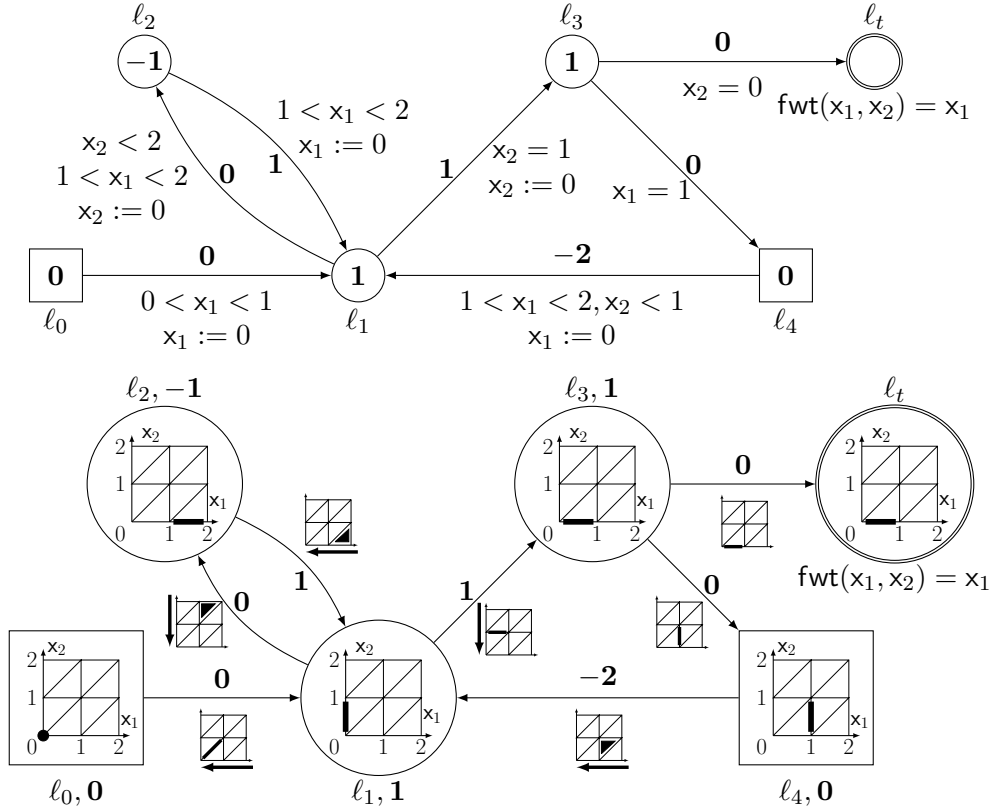


Figure 10: A weighted timed game \mathcal{G} with two clocks x_1 and x_2 , and the portion of its region game $\mathcal{R}(\mathcal{G})$ accessible from configuration $(l_0, (0,0))$. The states of $\mathcal{R}(\mathcal{G})$ are labelled by their associated region, location and weight, and edges are labelled by a representation of their guards and resets. For example, the edge from (l_0, r_0) to (l_1, r_1) in $\mathcal{R}(\mathcal{G})$ highlights the time successors of the region r_0 that satisfy the guard $0 < x_1 < 1$, and the arrow represents the direction in which this set of points is projected by the clock reset $x_1 := 0$, so that we end up in the region r_1 .

left of Figure 11): we must think about the final weight functions as the previously computed approximations of the values of SCCs coming after the current one in the topological order. We will keep as an invariant that final weight functions are piecewise affine with a finite number of pieces, and are continuous on each region.

For an SCC of $\mathcal{R}(\mathcal{G})$ and an initial state $(l_0, [\mathbf{0}])$ of $\mathcal{R}(\mathcal{G})$ provided by the SCC decomposition, we show that the game on the SCC is equivalent to a game on a tree built from a semi-unfolding (see middle of Figure 11) of $\mathcal{R}(\mathcal{G})$ from $(l_0, [\mathbf{0}])$ of finite depth, with certain nodes of the tree being *kernels* (parts of $\mathcal{R}(\mathcal{G})$ that contain all cycles of weight 0). The semi-unfolding is stopped either when reaching a final location, or when some location (or kernel) has been visited for a certain fixed number of times. Notice that, for divergent WTGs, there are no kernels, which simplifies the computation.

Then, we compute an approximation of $\text{Val}_{\mathcal{G}}(l_0, \mathbf{0})$ with a bottom-up computation on the semi-unfolding. This computation is exact on nodes labelled by a single region state s , but approximate on kernel nodes K_s . For the latter, we use the corner-point abstraction (right

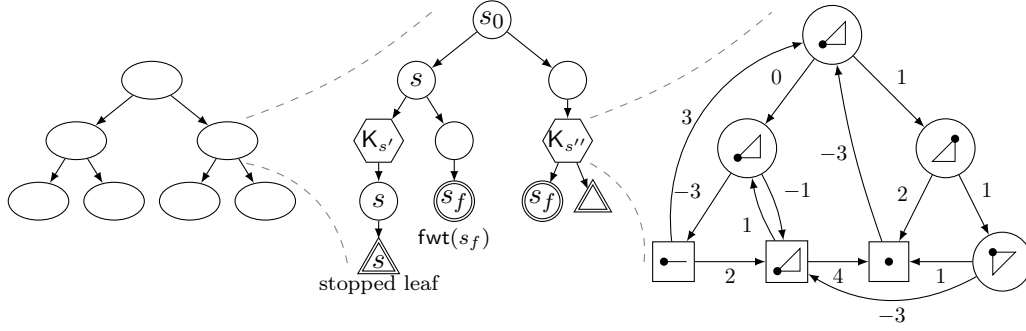


Figure 11: Static approximation schema: SCC decomposition of $\mathcal{R}(\mathcal{G})$, semi-unfolding of an SCC, corner-point abstraction for the kernels

of Figure 11) over $1/N$ -regions to compute values, and prove that, with an appropriately chosen N , this provides an ε -approximation of values.

This resolution of the value problem for divergent WTGs and approximation problem for almost-divergent WTGs heavily relies on the region abstraction, and requires one to construct $\mathcal{R}(\mathcal{G})$ entirely and compute its SCCs, before unfolding it partially in a tree-shaped structure. Our second result is a more symbolic approximation schema based on the value iteration algorithm only (in case we are able to rule out the presence of configurations with value $-\infty$, which could in particular be true if there are only non-negative weights): the computations are not performed on the region abstraction, but instead use the cell partitions introduced in Section 2.3 that can cover several regions.

Theorem 3.10. *Let \mathcal{G} be an almost-divergent WTG such that $\text{Val}_{\mathcal{G}}(\ell, \nu) > -\infty$ for every configuration (ℓ, ν) . Then the sequence $(\text{Val}_{\mathcal{G}}^k)_{k \geq 0}$ converges towards $\text{Val}_{\mathcal{G}}$ and for every $\varepsilon \in \mathbb{Q}_{>0}$, there exists an integer P such that $\text{Val}_{\mathcal{G}}^P$ is an ε -approximation of $\text{Val}_{\mathcal{G}}$ for all configurations.*

Note that we have to control for configurations (ℓ, ν) of value $-\infty$, where the non-increasing sequence $(\text{Val}_{\mathcal{G}}^k(\ell, \nu))_{k \in \mathbb{N}}$ (that starts at $+\infty$) will diverge towards $-\infty$, but has no hope of approximating it. However, we will show that the configurations with value $-\infty$ can be computed in advance:

Proposition 3.11. *Given an almost-divergent WTG \mathcal{G} and an initial location ℓ_0 , the decision problem asking whether $\text{Val}_{\mathcal{G}}(\ell_0, \mathbf{0}) = -\infty$ is EXPTIME-complete.*

The exponential upper bound is obtained in Section 5. This contrasts with the general case (not necessarily almost-divergent), where the $-\infty$ -value problem is undecidable (among other problems), as detailed in Proposition 3.12.

3.1. Hardness of value problems. The EXPTIME-hardness result of Theorem 3.4 (respectively Proposition 3.11) is a reduction from the problem of solving timed games with reachability objectives [JT07].

To a reachability timed game \mathcal{G} , we simply set the weight of each edge to 1 and the weight of each location to 0, making it a divergent WTG. We set the final weight of every target configuration at 0 (respectively $-\infty$). Then, Min wins the reachability timed game if and

only if the value in the WTG is lower than threshold $\alpha = |\mathcal{R}(\mathcal{G})|$ (respectively equals $-\infty$). One direction of this statement is immediate by definition of having a value smaller than $+\infty$, and the other comes from the fact that reachability in the timed game implies reachability in the region game in less than α transitions, in turn implying that Min can ensure target reachability in the WTG with cumulative weight below α , *i.e.* $\text{Val}_{\mathcal{G}}(s, \nu) \leq \alpha + \text{fwt}$, with $\text{fwt} = 0$ (respectively $-\infty$).

Proposition 3.12. *Given a WTG \mathcal{G} and an initial location ℓ_0 , the decision problem asking whether $\text{Val}_{\mathcal{G}}(\ell_0, \mathbf{0}) = -\infty$ is undecidable.*

Proof. The proof goes via a reduction to the existence problem on turn-based WTG: given a WTG \mathcal{G} (without final weight function), a non-negative integer threshold α and a starting location ℓ_0 , does there exist a strategy for Min that can guarantee reaching the unique target location ℓ_t from ℓ_0 with weight $< \alpha$. In the setting with non-negative weights in \mathcal{G} , it is proven in [BBM06] that the problem is undecidable for the comparison $\leq \alpha$. In the setting with arbitrary weights, formal proofs are given for all comparison signs in [BGNK⁺14].

Consider the WTG \mathcal{G}' built from \mathcal{G} by adding a transition from ℓ_t to ℓ_0 , without guards and resetting all the clocks, of discrete weight $-\alpha$. We add a new target location ℓ'_t , and add transitions of weight 0 from ℓ_t to ℓ'_t . Locations ℓ_t and ℓ'_t belong to Min. Let us prove that $\text{Val}_{\mathcal{G}'}(\ell_0, \mathbf{0}) = -\infty$ if and only if Min has a strategy to guarantee a weight $< \alpha$ in \mathcal{G} . Assume first $\text{Val}_{\mathcal{G}'}(\ell_0, \mathbf{0}) = -\infty$. If $\text{Val}_{\mathcal{G}}(\ell_0, \mathbf{0}) = -\infty$, we are done. Otherwise, Min must follow in \mathcal{G}' the new transition from ℓ_t to ℓ_0 to enforce a cycle of negative value, and thus enforce a play from $(\ell_0, \mathbf{0})$ to ℓ_t with weight less than α . Therefore, there exists a strategy for Min in \mathcal{G} that can guarantee a weight $< \alpha$. Reciprocally, if there exists a strategy for Min in \mathcal{G} that can guarantee a weight $< \alpha$, then Min can force a negative cycle play in \mathcal{G}' and $\text{Val}_{\mathcal{G}'}(\ell_0, \mathbf{0}) = -\infty$. \square

4. DECIDING DIVERGENCE AND ALMOST-DIVERGENCE

In this section, we will study properties that region cycles must satisfy in divergent or almost-divergent WTGs. This will give us a better understanding of the modelling power these classes confer, as well as enable us to provide procedures of optimal complexity to decide if a WTG fulfils the divergence or almost-divergence conditions.

4.1. Cycles in an almost-divergent WTG. Let us start with properties that hold for all almost-divergent weighted timed games \mathcal{G} . A region cycle \mathbf{p} of $\mathcal{R}(\mathcal{G})$ is said to be a positive cycle (respectively a negative cycle, a 0-cycle) if every finite play ρ following \mathbf{p} satisfies $\text{wt}(\rho) \geq 1$ (respectively $\text{wt}(\rho) \leq -1$, $\text{wt}(\rho) = 0$).

We start by showing that, in an almost-divergent game, all cycles $\mathbf{p} = t_1 \cdots t_k$ of $\mathcal{R}(\mathcal{G})$ (with t_1, \dots, t_k edges of $\mathcal{R}(\mathcal{G})$) are either 0-cycles, positive cycles or negative cycles¹⁰, and we can classify a cycle by looking only at one of the corner plays following it:

Lemma 4.1. *Let \mathcal{G} be an almost-divergent WTG. A region cycle \mathbf{p} is a positive cycle (respectively a negative cycle, a 0-cycle) if and only if there exists a corner play $\bar{\rho}$ following \mathbf{p} with $\text{wt}_{\Sigma}(\bar{\rho}) > 0$ (respectively $\text{wt}_{\Sigma}(\bar{\rho}) < 0$, $\text{wt}_{\Sigma}(\bar{\rho}) = 0$). Moreover, every region cycle in \mathcal{G} is either positive, negative, or a 0-cycle.*

¹⁰In contrast, Definition 3.6 only requires that each play following a region cycle has weight in $(-\infty, -1] \cup \{0\} \cup [1, +\infty)$, without disallowing a region cycle to contain plays of different types.

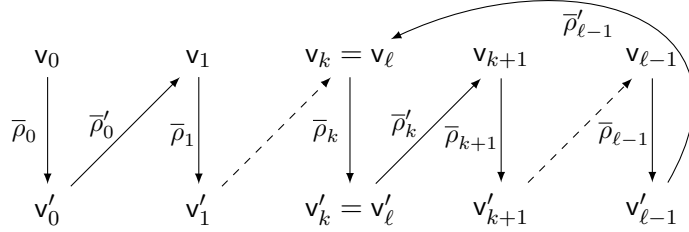


Figure 12: Proof scheme of Lemma 4.2. The top labels are corners of $\text{first}(\mathbf{p})$, the bottom ones are corners of $\text{first}(\mathbf{p}')$, and edges represent corner plays.

Proof. If \mathbf{p} is a positive cycle (respectively a negative cycle, a 0-cycle), every such corner play $\bar{\rho}$ will have weight above 0 (respectively under 0, equal to 0), by Lemma 2.16. Reciprocally, if such a corner play exists, all corner plays following \mathbf{p} have weight above 0 (respectively under 0, equal to 0): otherwise the set $\{\text{wt}_{\Sigma}(\rho) \mid \rho \text{ play following } \mathbf{p}\}$ would have non-empty intersection with the set $(-1, 0) \cup (0, 1)$ by Lemma 2.16, which would contradict that the game is almost-divergent.

Let \mathbf{p} be a region cycle of \mathcal{G} . All the plays following \mathbf{p} have a weight in $(-\infty, -1] \cup \{0\} \cup [1, +\infty)$. If it has a play in two different subintervals $(-\infty, -1]$, $\{0\}$, and $[1, +\infty)$, then Lemma 2.16 implies also that a play following \mathbf{p} will have a weight in $(-1, 0) \cup (0, 1)$, which is forbidden by Definition 3.6. \square

An important result is that the sign of cycles is stable by rotation. This is not trivial because plays following a cycle can start and end in different valuations, therefore changing the starting region state of the cycle could *a priori* change the plays that follow it and the sign of their weights.

Lemma 4.2. *Let \mathbf{p} and \mathbf{p}' be region paths of an almost-divergent WTG. If $\mathbf{p}\mathbf{p}'$ is a positive cycle (respectively a negative cycle, a 0-cycle), then $\mathbf{p}'\mathbf{p}$ is a positive cycle (respectively a negative cycle, a 0-cycle).*

Proof. Since $\mathbf{p}_1 = \mathbf{p}\mathbf{p}'$ is a cycle, $\text{first}(\mathbf{p}) = \text{last}(\mathbf{p}')$ and $\text{first}(\mathbf{p}') = \text{last}(\mathbf{p})$, so $\mathbf{p}_2 = \mathbf{p}'\mathbf{p}$ is a cycle as well. First, since there are finitely many corners, by constructing a long enough play following an iterate of $\mathbf{p}'\mathbf{p}$, we can obtain a corner play that starts and ends in the same corner. Formally, we define two sequences of region corners $(v_i \in \text{first}(\mathbf{p}))_i$ and $(v'_i \in \text{first}(\mathbf{p}'))_i$. We start by choosing any $v_0 \in \text{first}(\mathbf{p})$. Let v'_0 be a corner of $\text{first}(\mathbf{p}')$ such that v'_0 is accessible from v_0 by following \mathbf{p} with a corner play $\bar{\rho}_0$. For every $i > 0$, let v_i be a corner of $\text{first}(\mathbf{p})$ such that v_i is accessible from v'_{i-1} by following \mathbf{p}' with a corner play $\bar{\rho}'_{i-1}$, and let v'_i be a corner of $\text{first}(\mathbf{p}')$ such that v'_i is accessible from v_i by following \mathbf{p} with a corner play $\bar{\rho}_i$. We stop the construction at the first index ℓ such that there exists $k < \ell$ with $v_{\ell} = v_k$. Additionally, we let $\bar{\rho}_{\ell} = \bar{\rho}_k$. We know that this process never gets stuck—*i.e.* we can always find such corner plays iteratively—by Lemma 2.17, and it is bounded since $\text{first}(\mathbf{p})$ has at most $|\mathcal{X}| + 1$ corners.

For every $0 \leq i \leq \ell$, let w_i be the weight of the corner play $\bar{\rho}_i$ from v_i to v'_i along \mathbf{p} , and let w'_i be the weight of the corner play $\bar{\rho}'_i$ from v'_i to v_{i+1} along \mathbf{p}' . The concatenation of the two plays has weight $w_i + w'_i > 0$ (respectively $w_i + w'_i < 0$, $w_i + w'_i = 0$), since it follows the positive cycle (respectively negative cycle, 0-cycle) \mathbf{p}_1 . For every $0 \leq i < \ell$, the

concatenation of the corner play $\bar{\rho}'_i$ from v'_i to v_{i+1} with the corner play $\bar{\rho}_{i+1}$ from v_{i+1} to v'_{i+1} is a play from v'_i to v'_{i+1} , of weight $w'_i + w_{i+1}$, following \mathbf{p}_2 . Since \mathbf{p}_2 is a cycle, and the game is almost-divergent, all possible values of $w'_i + w_{i+1}$ have the same sign by Lemma 4.1.

Finally, we can construct a corner play from v'_k to v'_ℓ by concatenating the plays $\bar{\rho}'_k, \bar{\rho}_{k+1}, \dots, \bar{\rho}_{\ell-1}, \bar{\rho}'_{\ell-1}, \bar{\rho}_\ell$. We denote the weight of that play W , and

$$W = \sum_{i=k}^{\ell-1} (w'_i + w_{i+1}) = \sum_{i=k}^{\ell-1} (w_i + w'_i)$$

since $w_k = w_\ell$. As $w_i + w'_i > 0$ (respectively $w_i + w'_i < 0$, $w_i + w'_i = 0$) holds for every i , we obtain $W > 0$ (respectively $W < 0$, $W = 0$).

This implies that the terms $w'_i + w_{i+1}$, of constant sign, are all above 0 (respectively under 0, equal to 0). As a consequence, the concatenation of $\bar{\rho}'_k$ and $\bar{\rho}_{k+1}$ is a corner play following \mathbf{p}_2 of weight above 0 (respectively under 0, equal to 0). By Lemma 4.1, we conclude that \mathbf{p}_2 must be a positive cycle (respectively a negative cycle, a 0-cycle). \square

Therefore, region cycles in almost-divergent games are well-behaved: we can compose and rotate them while preserving their sign in the expected way.

4.2. SCC-based characterisations. After studying the properties of region cycles, we now focus on strongly connected components (SCCs) of the region abstraction $\mathcal{R}(\mathcal{G})$. An SCC S of $\mathcal{R}(\mathcal{G})$ is said to be positive (respectively negative) if every cycle in S is positive (respectively negative), *i.e.* if every play ρ following a region cycle in S satisfies $\text{wt}_\Sigma(\rho) \geq 1$ (respectively $\text{wt}_\Sigma(\rho) \leq -1$).

Proposition 4.3. *A weighted timed game \mathcal{G} is divergent if and only if, each SCC of $\mathcal{R}(\mathcal{G})$ is either positive or negative.*

Likewise, an SCC S of $\mathcal{R}(\mathcal{G})$ is said to be non-negative (respectively non-positive) if every region cycle in S is either a positive cycle or a 0-cycle (respectively either a negative cycle or a 0-cycle), *i.e.* every play ρ following a region cycle in S satisfies either $\text{wt}_\Sigma(\rho) \geq 1$ or $\text{wt}_\Sigma(\rho) = 0$ (respectively either $\text{wt}_\Sigma(\rho) \leq -1$ or $\text{wt}_\Sigma(\rho) = 0$). We obtain:

Proposition 4.4. *A WTG \mathcal{G} is almost-divergent if and only if each SCC of $\mathcal{R}(\mathcal{G})$ is either non-negative or non-positive.*

We now prove these two results. First, note that if \mathcal{G} is divergent it has no 0-cycle, and Proposition 4.4 implies that each SCC of $\mathcal{R}(\mathcal{G})$ is either positive or negative. Conversely, if each SCC of $\mathcal{R}(\mathcal{G})$ is either positive or negative, Proposition 4.4 implies that \mathcal{G} is divergent. Therefore, Proposition 4.3 is a corollary of Proposition 4.4. The rest of this section now proves Proposition 4.4.

To prove the reciprocal implication of Proposition 4.4, we only need to show that non-negative (respectively non-positive) SCCs of $\mathcal{R}(\mathcal{G})$ satisfy the definition of almost-divergent WTGs. By definition, they only contain executions ρ following region cycles \mathbf{p} such that $\text{wt}_\Sigma(\rho) \in \{0\} \cup [1, +\infty)$ (respectively $\text{wt}_\Sigma(\rho) \in (-\infty, -1] \cup \{0\}$). Then, assume $\text{wt}_\Sigma(\rho) = 0$ and that \mathbf{p} can be decomposed into smaller cycles \mathbf{p}' and \mathbf{p}'' . Definition 3.6 requires us to show that all plays ρ' and ρ'' following \mathbf{p}' and \mathbf{p}'' , respectively, are such that $\text{wt}_\Sigma(\rho') = \text{wt}_\Sigma(\rho'') = 0$, *i.e.* \mathbf{p}' and \mathbf{p}'' are 0-cycles. Note that by Lemma 4.2, $\mathbf{p}'\mathbf{p}''$ is a 0-cycle. As \mathbf{p}' and \mathbf{p}'' are contained in the same SCC, they are either both non-negative cycles or both non-positive cycles. Let $\rho'\rho''$ be a play following $\mathbf{p}'\mathbf{p}''$, so that ρ' follows \mathbf{p}' and ρ'' follows \mathbf{p}'' . Then,

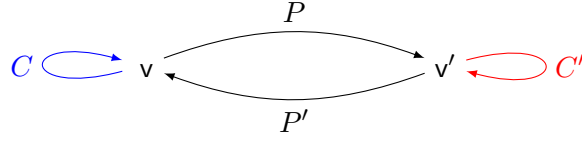


Figure 13: Proof scheme of Lemma 4.5, with paths and cycles around corners of the graph $\text{FOG}(\mathbf{p}, \mathbf{p}')$.

$\text{wt}_\Sigma(\rho' \rho'') = \text{wt}_\Sigma(\rho') + \text{wt}_\Sigma(\rho'') = 0$, it follows that $\text{wt}_\Sigma(\rho') = \text{wt}_\Sigma(\rho'') = 0$, *i.e.* \mathbf{p}' and \mathbf{p}'' are 0-cycles.

For the direct implication, the situation is more complex: we need to be careful while composing cycles with each other. To help us, we rely on the folded orbit graphs of region cycles. Suppose that \mathcal{G} is almost-divergent, and consider two cycles \mathbf{p} and \mathbf{p}' in the same SCC of $\mathcal{R}(\mathcal{G})$. We need to show that they are both either non-positive or non-negative. Lemma 4.5 will first take care of the case where \mathbf{p} and \mathbf{p}' share a region state (ℓ, r) .

Lemma 4.5. *If \mathcal{G} is almost-divergent and two cycles \mathbf{p} and \mathbf{p}' of $\mathcal{R}(\mathcal{G})$ share a region state (ℓ, r) , they are either both non-negative or both non-positive.*

Proof. Suppose by contradiction that \mathbf{p} is negative and \mathbf{p}' is positive. We assume that (ℓ, r) is the first region state of both \mathbf{p} and \mathbf{p}' , possibly performing a rotation of the cycles if necessary (in particular this preserves their sign by Lemma 4.2). We will rely on the folded orbit graphs that have been defined on page 19. We construct a graph $\text{FOG}(\mathbf{p}, \mathbf{p}')$ as the union of $\text{FOG}(\mathbf{p})$ and $\text{FOG}(\mathbf{p}')$ (that share the same set of vertices), colouring in blue the edges of $\text{FOG}(\mathbf{p})$ and in red the edges of $\text{FOG}(\mathbf{p}')$. A path in $\text{FOG}(\mathbf{p}, \mathbf{p}')$ is said blue (respectively red) when all of its edges are blue (respectively red).

Since $\text{FOG}(\mathbf{p})$ and $\text{FOG}(\mathbf{p}')$ are finite graphs with no deadlocks (every corner has an outgoing edge by Lemma 2.17), from every corner of $\text{FOG}(\mathbf{p}, \mathbf{p}')$, we can reach a blue simple cycle, as well as a red simple cycle. Since there are only a finite number of simple cycles in $\text{FOG}(\mathbf{p}, \mathbf{p}')$, there exists a blue cycle C and a red cycle C' that can reach each other in $\text{FOG}(\mathbf{p}, \mathbf{p}')$. Denote by \mathbf{v} and \mathbf{v}' the first corners of cycles C and C' , respectively.

We assume first $\mathbf{v} = \mathbf{v}'$. Let k and k' be the respective lengths of C and C' , so that C can be decomposed as $\mathbf{v} \xrightarrow{\bar{\rho}_1} \dots \xrightarrow{\bar{\rho}_k} \mathbf{v}$ and C' as $\mathbf{v} \xrightarrow{\bar{\rho}'_1} \dots \xrightarrow{\bar{\rho}'_{k'}} \mathbf{v}$, where $\bar{\rho}_i$ are corner plays following \mathbf{p} and $\bar{\rho}'_i$ are corner plays following \mathbf{p}' . Let $\bar{\rho}$ be the concatenation of $\bar{\rho}_1, \dots, \bar{\rho}_k$, and $\bar{\rho}'$ be the concatenation of $\bar{\rho}'_1, \dots, \bar{\rho}'_{k'}$. Recall that $w = |\text{wt}_\Sigma(\bar{\rho})|$ and $w' = |\text{wt}_\Sigma(\bar{\rho}')|$ are integers. Since \mathbf{p} is negative, so is \mathbf{p}^k , the concatenation of k copies of \mathbf{p} (the weight of a play following it is a sum of weights all below -1). Therefore, $\bar{\rho}$, that follows \mathbf{p}^k , has a weight $\text{wt}_\Sigma(\bar{\rho}) \leq -1$ by Lemma 2.16. Similarly, $\text{wt}_\Sigma(\bar{\rho}') \geq 1$. Let $\bar{\rho}''$ be the play obtained by concatenating w' copies of $\bar{\rho}$ and w copies of $\bar{\rho}'$. Then, $\text{wt}_\Sigma(\bar{\rho}'') = \text{wt}_\Sigma(\bar{\rho})w' + \text{wt}_\Sigma(\bar{\rho}')w = 0$, and therefore the region cycle \mathbf{p}'' composed of w' copies of \mathbf{p}^k and w copies of $\mathbf{p}'^{k'}$ is a 0-cycle. This contradicts the almost-divergence of \mathcal{G} , since \mathbf{p}'' can be decomposed into smaller cycles that are not 0-cycles.

Therefore, \mathbf{v} and \mathbf{v}' are different, but can reach each other in $\text{FOG}(\mathbf{p}, \mathbf{p}')$. Let P be a path from \mathbf{v} to \mathbf{v}' , and P' be a path from \mathbf{v}' to \mathbf{v} . The situation is depicted in Figure 13. Consider the cycle C'' obtained by concatenating P and P' . As a cycle of $\text{FOG}(\mathbf{p}, \mathbf{p}')$, we can map it to a cycle \mathbf{p}'' of $\mathcal{R}(\mathcal{G})$ (alternating \mathbf{p} and \mathbf{p}' depending on the colours of the traversed edges), so that C'' is a cycle (of length 1) of $\text{FOG}(\mathbf{p}'')$. As \mathcal{G} is almost-divergent, \mathbf{p}'' is either

positive, negative or a 0-cycle. Moreover, \mathbf{p}'' cannot be a 0-cycle as it can be decomposed into smaller cycles that are not 0-cycles. Suppose for instance that it is positive. Since (ℓ, r) is the first region state of both \mathbf{p} and \mathbf{p}'' , we can construct the FOG(\mathbf{p}, \mathbf{p}''), in which C is a blue cycle and C'' is a red cycle, both sharing the same first vertex. We then conclude with the previous case. A similar reasoning with \mathbf{p}' applies to the case that \mathbf{p}'' is negative. Therefore, in all cases, we reach a contradiction. \square

To finish the proof of the direct implication of Proposition 4.4, we suppose that the two cycles \mathbf{p} and \mathbf{p}' , one positive and the other negative, in the same SCC of $\mathcal{R}(\mathcal{G})$, do not share any region states. By strong connectivity, in $\mathcal{R}(\mathcal{G})$, there exists a path \mathbf{p}_1 from the first state of \mathbf{p} to the first state of \mathbf{p}' , and a path \mathbf{p}_2 from the first state of \mathbf{p}' to the first state of \mathbf{p} . Consider the cycle \mathbf{p}'' of $\mathcal{R}(\mathcal{G})$ defined by $\mathbf{p}\mathbf{p}_1\mathbf{p}'\mathbf{p}_2$. By the almost-divergence of \mathcal{G} , \mathbf{p}'' must be either positive, negative or a 0-cycle. Since it shares a state with both \mathbf{p} and \mathbf{p}' , Lemma 4.5 allows us to prove a contradiction in both positive and negative cases, and therefore \mathbf{p}'' must be a 0-cycle. This contradicts the hypothesis as one of the decompositions of \mathbf{p}'' into smaller cycles produces \mathbf{p} and $\mathbf{p}_1\mathbf{p}'\mathbf{p}_2$, with \mathbf{p} a non-0-cycle. This concludes the proof of Proposition 4.4.

Remark 4.6. These characterisations of divergent or almost-divergent WTGs in term of SCCs provide an intuitive understanding of the modelling power these classes hold. For divergence, the model should have a global structure (the SCC decomposition) linking modules in an acyclic fashion. For each module, we have to choose between a positive dynamic, where weights always eventually increase, and a negative dynamic, where weights always eventually decrease. For almost-divergence, the modules may also have portions that are (eventually) neutral with regard to weight accumulation. In both classes, arbitrarily small weights should not be allowed to accumulate.

4.3. Deciding membership. We study the *membership problem* for divergent (respectively almost-divergent) WTGs, *i.e.* the decision problem that asks if a given WTG is divergent (respectively almost-divergent). As mentioned in Theorems 3.4 and 3.9, we show that it is PSPACE-complete for both of these classes.

Relying on the previous characterisation of Propositions 4.3 and 4.4, the algorithms will consist in only considering region cycles of length bounded by the number of corners in the corner-point abstraction $\Gamma(\mathcal{G})$. For divergent WTGs, this will be correct by using the following result:

Lemma 4.7. *Let \mathcal{G} be a WTG. An SCC S of $\mathcal{R}(\mathcal{G})$ is positive (respectively negative) if and only if every region cycle in S , of length at most $|\Gamma(\mathcal{G})|$, is positive (respectively negative).*

Proof. The direct implication holds by definition. Reciprocally, let us assume that every cycle in S of length at most $|\Gamma(\mathcal{G})|$ is positive (respectively negative), and prove that every cycle \mathbf{p} in S is positive (respectively negative), by induction on the length of \mathbf{p} . If \mathbf{p} has length above $|\Gamma(\mathcal{G})|$, every corner play $\bar{\rho}$ following \mathbf{p} can be split as $\bar{\rho} = \bar{\rho}_1\bar{\rho}_2\bar{\rho}_3$, with $\bar{\rho}_2$ a corner play that starts and ends in the same corner. Then we can write $\mathbf{p} = \mathbf{p}_1\mathbf{p}_2\mathbf{p}_3$, with $\bar{\rho}_1$ (respectively $\bar{\rho}_2, \bar{\rho}_3$) following \mathbf{p}_1 (respectively $\mathbf{p}_2, \mathbf{p}_3$). Observe that \mathbf{p}_2 and $\mathbf{p}_1\mathbf{p}_3$ are region cycles of S , both positive (respectively negative) by induction. It follows that $\text{wt}_\Sigma(\bar{\rho}_2) \geq 1$ (respectively $\text{wt}_\Sigma(\bar{\rho}_2) \leq -1$), and $\text{wt}_\Sigma(\bar{\rho}_1\bar{\rho}_3) \geq 1$ (respectively $\text{wt}_\Sigma(\bar{\rho}_1\bar{\rho}_3) \leq -1$), as $\bar{\rho}_1\bar{\rho}_3$ is a valid corner play following $\mathbf{p}_1\mathbf{p}_3$. We can therefore conclude that $\text{wt}_\Sigma(\bar{\rho}) \geq 1$

(respectively $\text{wt}_\Sigma(\bar{\rho}) \leq -1$). This holds for all corner plays following \mathbf{p} , and by Lemma 2.16 \mathbf{p} is positive (respectively negative). \square

Then, to decide if a game is *not divergent*, using Proposition 4.3, it suffices to search for an SCC of the region automaton containing a cycle such that there exists a corner play following it of non-negative weight, and a cycle such that there exists a corner play following it of non-positive weight, both of length bounded by $B = |\Gamma(\mathcal{G})| \leq |L| \times |\text{Reg}(\mathcal{X}, M)| \times (|\mathcal{X}| + 1)$. We can test this condition in NPSPACE: we guess a starting region for each cycle, use standard reachability analysis [AD94] to check that they are in the same SCC of $\mathcal{R}(\mathcal{G})$ (in PSPACE), and use the following result with comparison ≥ 0 and ≤ 0 , respectively, to check the sign of each cycle.

Lemma 4.8. *Consider a WTG \mathcal{G} , a region state (ℓ, r) of $\mathcal{R}(\mathcal{G})$, a bound $B \in \mathbb{N}$, and a comparison operator $\bowtie \in \{<, >, \leq, \geq, =, \neq\}$. Deciding if there exists a corner play $\bar{\rho}$ following a cycle \mathbf{p} of $\mathcal{R}(\mathcal{G})$ starting from (ℓ, r) , such that $|\mathbf{p}| \leq B$ and $\text{wt}_\Sigma(\bar{\rho}) \bowtie 0$, is in PSPACE (i.e. it can be done using space polynomial in $|\mathcal{G}|$ and $\log(B)$).*

Proof. We guess a starting corner v of r for $\bar{\rho}$, and we guess on-the-fly the transitions of \mathbf{p} and $\bar{\rho}$, i.e. a sequences of regions with one of their corners, keeping in memory the cumulative weight of $\bar{\rho}$ and the length $|\mathbf{p}|$. At every step, we check that $|\mathbf{p}| \leq B$ in space polynomial in $\log(B)$ and $\log(|\mathbf{p}|) \leq \log(|\mathcal{R}(\mathcal{G})|)$, with $\log(|\mathcal{R}(\mathcal{G})|)$ polynomial in $|\mathcal{G}|$.¹¹ Similarly, we can check that $\bar{\rho}$ is following \mathbf{p} in polynomial space. At some point, we guess that the cycle is complete, and we check that the current region state equals (ℓ, r) . Finally, we check that $\text{wt}_\Sigma(\bar{\rho}) \bowtie 0$ in space polynomial in $\log(\text{wt}_\Sigma(\bar{\rho}))$. Note that $\text{wt}_\Sigma(\bar{\rho})$ is an integer bounded (in absolute value) by $B \times w_{\max}$, and can thus be stored in polynomial space. This shows that the problem is in NPSPACE, and thus in PSPACE using Savitch's theorem [Sav70]. \square

Since the bound B is at most exponential in $|\mathcal{G}|$, this check can be performed in PSPACE. This shows that the membership problem for divergent weighted timed games is in $\text{coNPSPACE} = \text{coPSPACE} = \text{PSPACE}$ by Savitch [Sav70].

Let us now show the (co) PSPACE-hardness by a reduction from the reachability problem in a timed automaton. Consider a timed automaton with a starting location and a different target location without outgoing edges. We construct from it a weighted timed game by distributing all locations to Min, and equipping all edges with weight 1, and all locations with weight 0. We also add a loop with weight -1 on the target, and an edge from the target location to the initial location with weight 0, both with guard \top and resetting all clocks. Then, the WTG is not divergent if and only if the target can be reached from the initial location in the timed automaton.

For almost-divergent WTGs, the length of the required region cycles is bigger, because of the possible presence of 0-cycles.

Lemma 4.9. *Let \mathcal{G} be a WTG. An SCC S of $\mathcal{R}(\mathcal{G})$ is non-negative (respectively non-positive) if and only if every region cycle in S , of length at most $|\Gamma(\mathcal{G})|^2$, is either a positive cycle or a 0-cycle (respectively either a negative cycle or a 0-cycle).*

Proof. The direct implication holds by definition. Reciprocally, suppose that every cycle in S of length at most $|\Gamma(\mathcal{G})|^2$ is either a positive cycle or a 0-cycle (respectively either a negative cycle or a 0-cycle). Let us prove that every cycle \mathbf{p} in S is either a positive cycle

¹¹The global clock bound M is at most exponential in the size of \mathcal{G} , and $|\mathcal{R}(\mathcal{G})|$ is at most exponential in $|\mathcal{X}|$ but polynomial in M , therefore $|\mathcal{R}(\mathcal{G})|$ is at most exponential in $|\mathcal{G}|$.

or a 0-cycle (respectively either a negative cycle or a 0-cycle), by induction on the length of \mathbf{p} . Consider a region cycle \mathbf{p} with length above $|\Gamma(\mathcal{G})|^2$. Let us show that for all corner plays $\bar{\rho}$, $\bar{\rho}'$ following \mathbf{p} , either $\text{wt}_\Sigma(\bar{\rho}) = \text{wt}_\Sigma(\bar{\rho}') = 0$, or both $\text{wt}_\Sigma(\bar{\rho}) \geq 1$ and $\text{wt}_\Sigma(\bar{\rho}') \geq 1$ (respectively both $\text{wt}_\Sigma(\bar{\rho}) \leq -1$ and $\text{wt}_\Sigma(\bar{\rho}') \leq -1$) hold. This will allow us to conclude by Lemma 2.16.

From a pair of corner plays $\bar{\rho}$ and $\bar{\rho}'$ following \mathbf{p} , we can extract a sequence of pairs of corners states $((\ell_i, r_i, \mathbf{v}_i), (\ell_i, r_i, \mathbf{v}'_i))$, such that (ℓ_i, r_i) is the i -th region state of \mathbf{p} , and \mathbf{v}_i (respectively \mathbf{v}'_i) is the i -th corner of $\bar{\rho}$ (respectively $\bar{\rho}'$). Since $|\mathbf{p}| > |\Gamma(\mathcal{G})|^2$, there must exist two indexes, j and k , such that $j < k$, $(\ell_j, r_j) = (\ell_k, r_k)$ and $(\mathbf{v}_j, \mathbf{v}'_j) = (\mathbf{v}_k, \mathbf{v}'_k)$. In other words, we can write $\mathbf{p} = \mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3$, with \mathbf{p}_2 and $\mathbf{p}_1 \mathbf{p}_3$ region cycles of S , and $\bar{\rho}$, $\bar{\rho}'$ can be split as $\bar{\rho} = \bar{\rho}_1 \bar{\rho}_2 \bar{\rho}_3$, $\bar{\rho}' = \bar{\rho}'_1 \bar{\rho}'_2 \bar{\rho}'_3$, with $\bar{\rho}_l$ (respectively $\bar{\rho}'_l$) following \mathbf{p}_l (respectively \mathbf{p}'_l) for $l \in \{1, 2, 3\}$, such that $\bar{\rho}_2$ and $\bar{\rho}'_2$ are corner cycles, *i.e.* $\text{first}(\bar{\rho}_2) = \text{last}(\bar{\rho}_2)$ and $\text{first}(\bar{\rho}'_2) = \text{last}(\bar{\rho}'_2)$. Then, by induction either $\text{wt}_\Sigma(\bar{\rho}_2) = \text{wt}_\Sigma(\bar{\rho}'_2) = 0$, or both $\text{wt}_\Sigma(\bar{\rho}_2) \geq 1$ and $\text{wt}_\Sigma(\bar{\rho}'_2) \geq 1$ (respectively both $\text{wt}_\Sigma(\bar{\rho}_2) \leq -1$ and $\text{wt}_\Sigma(\bar{\rho}'_2) \leq -1$) hold. The same property holds for $\bar{\rho}_1 \bar{\rho}_3$ and $\bar{\rho}'_1 \bar{\rho}'_3$, both valid corner plays following $\mathbf{p}_1 \mathbf{p}_3$. It follows that either $\text{wt}_\Sigma(\bar{\rho}) = \text{wt}_\Sigma(\bar{\rho}') = 0$, or both $\text{wt}_\Sigma(\bar{\rho}) \geq 1$ and $\text{wt}_\Sigma(\bar{\rho}') \geq 1$ (respectively both $\text{wt}_\Sigma(\bar{\rho}) \leq -1$ and $\text{wt}_\Sigma(\bar{\rho}') \leq -1$) hold. \square

Then, to decide if a game is *not almost-divergent*, we distinguish two cases:

- There exists a region cycle, of length at most $B = |\Gamma(\mathcal{G})|^2$, and two corner plays $\bar{\rho}$ and $\bar{\rho}'$, both following \mathbf{p} , such that $\text{wt}_\Sigma(\bar{\rho}) = 0$ and $\text{wt}_\Sigma(\bar{\rho}') \neq 0$.
- An SCC of the region automaton contains a cycle such that there exists a corner play following it of negative weight, and a cycle such that there exists a corner play following it of positive weight, both of length bounded by $B = |\Gamma(\mathcal{G})|^2$.

We can test both conditions in **NPSPACE**, by guessing the starting regions of these cycles and using respectively Lemma 4.8 (for the second condition) and the following result (for the first condition):

Lemma 4.10. *Consider a weighted timed game \mathcal{G} , a region state (ℓ, r) of $\mathcal{R}(\mathcal{G})$, a bound $B \in \mathbb{N}$, and comparison operators $\bowtie, \bowtie' \in \{<, >, \leq, \geq, =, \neq\}$. Deciding if there exists a cycle \mathbf{p} of $\mathcal{R}(\mathcal{G})$ starting from (ℓ, r) , and two corner plays $\bar{\rho}$ and $\bar{\rho}'$, both following \mathbf{p} , such that $|\mathbf{p}| \leq B$, $\text{wt}_\Sigma(\bar{\rho}) \bowtie 0$ and $\text{wt}_\Sigma(\bar{\rho}') \bowtie' 0$, is in **PSPACE**.*

Proof. We follow the same non-deterministic procedure as Lemma 4.8, but this time we guess two corner plays on-the-fly instead of one. \square

This shows that the membership problem for divergent weighted timed games is in **coNPSPACE** = **coPSPACE** = **PSPACE** [Sav70].

Let us now show the (co) **PSPACE**-hardness by a reduction from the reachability problem in a timed automaton, similar to the one we used for the **PSPACE**-hardness of deciding divergence. We consider a timed automaton with a starting location and a different target location without outgoing edges. We construct from it a weighted timed game by distributing all locations to **Min**, and equipping all edges with weight 0, and all locations with weight 0. We also add a loop with weight 1 on the initial location, one with weight -1 on the target location, and an edge from the target location to the initial location with weight 0, all three with guard \top and resetting all clocks. Then, the weighted timed game is not almost-divergent if and only if the target can be reached from the initial location in the timed automaton.

5. DECIDING INFINITE VALUES

There are three reasons for an infinite value to appear in a WTG:

- an infinite value ($+\infty$ or $-\infty$) could appear in a final weight function and be propagated along a play;
- a value $+\infty$ could be obtained if Min is not able to reach a target location;
- a value $-\infty$ could be obtained if Min is able to reach a negative cycle, loop arbitrarily many times inside, and then reach a target location.

This section explains how to detect the three situations in the region game $\mathcal{R}(\mathcal{G})$ of an almost-divergent WTG \mathcal{G} . Before that, let us start by formalising a way to safely remove region states of $\mathcal{R}(\mathcal{G})$ for which the value of all their associated configurations is known to be infinite. Let $S^{+\infty}$ be a subset of $S = L \times \text{Reg}(\mathcal{X}, M)$, such that $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu) = +\infty$ for all $(\ell, r) \in S^{+\infty}$ and $\nu \in r$. If a configuration of \mathcal{G} is in the attractor of Max towards $S^{+\infty}$ (see Remark 2.18), then Max has a strategy giving it value $+\infty$. Moreover, the attractor being computable using regions, either all configurations in the same region have value $+\infty$, or none have value $+\infty$: therefore we could add this region to $S^{+\infty}$. We will thus assume that $S^{+\infty}$ is *closed by attractor* for Max, *i.e.* the attractor of Max towards $S^{+\infty}$ equals $S^{+\infty}$. We can define the same notion for a set $S^{-\infty}$ of states with value $-\infty$, that can be assumed closed by attractor of Min.

Lemma 5.1. *In $\mathcal{R}(\mathcal{G})$, let $S^{+\infty}$ be a set of region states of value $+\infty$ closed by attractor of Max, and let $S^{-\infty}$ be a set of region states of value $-\infty$ closed by attractor of Min. Removing the region states $S^{+\infty} \cup S^{-\infty}$ from $\mathcal{R}(\mathcal{G})$ will not change any other value.*

Proof. We only prove the result for the removal of region states from $S^{-\infty}$, since the proof for $S^{+\infty}$ is entirely symmetrical. Let $S' = S'_{\text{Min}} \uplus S'_{\text{Max}}$ denote $S \setminus S^{-\infty}$, and let $\mathcal{R}(\mathcal{G})'$ denote the restriction of $\mathcal{R}(\mathcal{G})$ to S' . Let us show that for all region states $(\ell, r) \in S'$ and $\nu \in r$, $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu) = \text{Val}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu)$.

If σ_{Min} is a strategy of Min in $\mathcal{R}(\mathcal{G})$, and σ'_{Min} is a strategy of Min in $\mathcal{R}(\mathcal{G})'$, such that for all plays ρ in $\text{FPlays}_{\mathcal{R}(\mathcal{G})'}^{\text{Min}}$, $\sigma_{\text{Min}}(\rho) = \sigma'_{\text{Min}}(\rho)$, we say that σ_{Min} coincides with σ'_{Min} . We define the same notion for strategies of Max. All strategies of Min and Max in $\mathcal{R}(\mathcal{G})'$ can be extended arbitrarily to become strategies in $\mathcal{R}(\mathcal{G})$ that coincide, by making the same choices on plays that stay in $\mathcal{R}(\mathcal{G})'$, and making arbitrary choices otherwise. It follows that if σ_{Min} and σ'_{Min} are strategies of Min that coincide, then for all strategies σ'_{Max} of Max in $\mathcal{R}(\mathcal{G})'$, there exists a corresponding strategy σ_{Max} in $\mathcal{R}(\mathcal{G})$, such that $\text{play}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu, \sigma'_{\text{Min}}, \sigma'_{\text{Max}}) = \text{play}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu, \sigma_{\text{Min}}, \sigma_{\text{Max}})$, and therefore

$$\text{Val}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu, \sigma'_{\text{Min}}) \leq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu, \sigma_{\text{Min}})$$

as the supremum over strategies of Max in $\mathcal{R}(\mathcal{G})$ is at least equal to the one in $\mathcal{R}(\mathcal{G})'$. Similarly, if σ_{Max} and σ'_{Max} are strategies of Max that coincide, then

$$\text{Val}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu, \sigma'_{\text{Max}}) \geq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu, \sigma_{\text{Max}})$$

By closure by attractor of Min, all transitions starting in a region state of S'_{Min} must end in a region state of S' (otherwise the first state would belong to the attractor of Min towards $S^{-\infty}$). Therefore, every strategy of Min in $\mathcal{R}(\mathcal{G})$ has a corresponding strategy in $\mathcal{R}(\mathcal{G})'$ that makes the same choices. In particular, if we consider an ε -optimal strategy σ_{Min} of Min in $\mathcal{R}(\mathcal{G})$, and its corresponding strategy σ'_{Min} in $\mathcal{R}(\mathcal{G})'$, it holds that

$$\text{Val}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu) \leq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu, \sigma'_{\text{Min}})$$

by definition of Val as an infimum over strategies, and

$$\text{Val}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu), \sigma'_{\text{Min}}) \leq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu), \sigma_{\text{Min}})$$

as explained before. Finally, $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu), \sigma_{\text{Min}}) \leq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu) + \varepsilon$ by ε -optimality of σ_{Min} , and thus $\text{Val}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu) \leq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu) + \varepsilon$. Since this holds for all $\varepsilon > 0$, we get that $\text{Val}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu) \leq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu)$.

From every state s of S'_{Max} , a positional strategy that chooses a transition jumping into $S^{-\infty}$ must have value $-\infty$. Any other choice would thus be equal or better, and by closure by attractor of Min , there must exist a transition starting in (ℓ, r) that ends in S' . Therefore, there exists an ε -optimal strategy σ_{Max} of Max in $\mathcal{R}(\mathcal{G})$ with a corresponding strategy σ'_{Max} in $\mathcal{R}(\mathcal{G})'$. Then, it holds that $\text{Val}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu) \geq \text{Val}_{\mathcal{R}(\mathcal{G})'}((\ell, r), \nu), \sigma'_{\text{Max}}) \geq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu), \sigma_{\text{Max}}) \geq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu) - \varepsilon$. We conclude since this holds for all ε . \square

5.1. Infinite final values. As a first step, we explain how to compute and remove all region states with value $+\infty$, and region states with value $-\infty$ because of final weights. The only states with infinite value that will remain are some states that derive a value of $-\infty$ from arbitrary accumulation of negative weights, and we will deal with them later.

Recall that the final weight function fwt has been supposed piecewise affine with a finite number of pieces and is continuous on each region. In particular, final weights $+\infty$ or $-\infty$ are given to entire regions. Then, let $S_t^{-\infty} \subseteq L_t \times \text{Reg}(\mathcal{X}, M)$ (respectively $S_t^{+\infty}$) denote the set of target region states that fwt maps to $-\infty$ (respectively $+\infty$).

Proposition 5.2. *If a region state of \mathcal{G} is in the attractor of Min towards $S_t^{-\infty}$ (respectively in the attractor of Max towards $S_t^{+\infty}$), then it has value $-\infty$ (respectively $+\infty$). Moreover, if we remove those states from $\mathcal{R}(\mathcal{G})$, the value of the other configurations does not change.*

Proof. If a region state (ℓ, r) is in the attractor of Min towards $S_t^{-\infty}$, then clearly Min has a (positional) strategy giving value $-\infty$ from (ℓ, r) , and $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu) = -\infty$ for all $\nu \in r$. An attractor of Min is always closed by attractor of Min , so we can apply Lemma 5.1 and conclude. Once again, a symmetrical proof lets us deal with the attractor of Max towards $S_t^{+\infty}$. \square

Then, assuming that all final weights are finite, configurations with value $+\infty$ are those from which Min cannot reach the target region states: thus, they can also be computed and removed using the attractor algorithm.

Proposition 5.3. *If fwt maps all configurations to values in \mathbb{R} , then a configuration $((\ell, r), \nu)$ has value $+\infty$ if and only if (ℓ, r) is not in the attractor of Min towards region states $L_t \times \text{Reg}(\mathcal{X}, M)$. Moreover, if we remove those region states from $\mathcal{R}(\mathcal{G})$, the value of the other configuration does not change.*

Proof. If a region state (ℓ, r) is not in the attractor of Min towards $L_t \times \text{Reg}(\mathcal{X}, M)$, then Max has a (safety) strategy giving value $+\infty$ from all configurations $((\ell, r), \nu)$ with $\nu \in r$, and $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu) = +\infty$. Conversely, if (ℓ, r) is in the attractor of Min towards $L_t \times \text{Reg}(\mathcal{X}, M)$, then Min has a strategy giving finite value from $((\ell, r), \nu)$, and $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell, r), \nu) < +\infty$. We conclude by Lemma 5.1. \square

We can now assume that all configurations have value in $\mathbb{R} \cup \{-\infty\}$ and that all target region states have final weight in \mathbb{R} : the precomputation needed so far consists only of attractor computations, that can be performed in time linear in $|\mathcal{R}(\mathcal{G})|$.

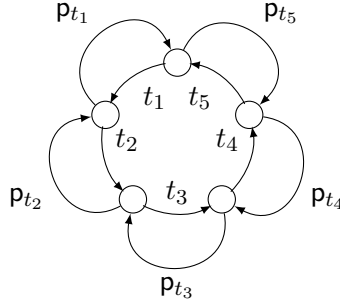


Figure 14: Application of the claim in the proof of Proposition 5.4

As previously explained, finding all states of value $-\infty$ is harder (in particular undecidable in general), but whenever we do manage to find them we can safely remove them by Lemma 5.1. The solution is quite simple for divergent WTGs, and slightly more involved in almost-divergent WTGs, because of the appearance of 0-cycles. In order to only present one uniform solution, we directly deal with the more general almost-divergent case. To do so, we introduce a new tool, the *kernels*, that will also be very useful in the rest of the study.

5.2. Kernel of an almost-divergent WTG. Kernels of an almost-divergent WTG are a way to group all 0-cycles of the game. We study those kernels and give a characterisation allowing computability. In [BJM15], kernels have also been studied, and contain exactly all transitions and locations of weight 0. Contrary to their non-negative case, the situation is more complex in our case with arbitrary weights since 0-cycles could go through locations or transitions that have weight different from 0. Moreover, it is not trivial (and may not be true in a non almost-divergent WTG) to know whether it is sufficient to consider only simple 0-cycles, *i.e.* cycles without repetitions.

We will now construct the kernel K as the subgraph of $\mathcal{R}(\mathcal{G})$ containing all 0-cycles. Formally, let T_K be the set of edges of $\mathcal{R}(\mathcal{G})$ belonging to a *simple* 0-cycle, and S_K be the set of states covered by T_K . We define the kernel K of $\mathcal{R}(\mathcal{G})$ as the subgraph of $\mathcal{R}(\mathcal{G})$ defined by S_K and T_K . Edges in $T \setminus T_K$ with starting state in S_K are called the output edges of K . We define it using only simple 0-cycles in order to ensure its computability. However, we now show that this is of no harm, since the kernel contains exactly all the 0-cycles, which will be crucial in our approximation schema.

Proposition 5.4. *A cycle of $\mathcal{R}(\mathcal{G})$ is entirely in K if and only if it is a 0-cycle.*

Proof. We prove that every 0-cycle is in K by induction on the length of the cycles. The initialisation contains only cycles of length 1, that are in K by construction. If we consider a cycle \mathbf{p} of length above 1, it is either simple (and thus in K , by definition), or it can be rotated and decomposed into $\mathbf{p}'\mathbf{p}''$, \mathbf{p}' and \mathbf{p}'' being smaller cycles. Let ρ be a corner play following $\mathbf{p}'\mathbf{p}''$. We denote by ρ' the prefix of ρ following \mathbf{p}' and ρ'' the suffix following \mathbf{p}'' . It holds that $\text{wt}_\Sigma(\rho') = -\text{wt}_\Sigma(\rho'')$, and in an almost-divergent SCC this implies $\text{wt}_\Sigma(\rho') = \text{wt}_\Sigma(\rho'') = 0$. Therefore, by Lemma 4.1 both \mathbf{p}' and \mathbf{p}'' are 0-cycles, and they must be in K by induction hypothesis.

We now prove that every cycle in K is a 0-cycle. By construction, every edge $t \in T_K$ is part of a simple 0-cycle. Thus, to every edge $t \in T_K$, we can associate a path \mathbf{p}_t such that

$t\mathbf{p}_t$ is a simple 0-cycle (rotate the simple cycle if necessary). The situation is exemplified in Figure 14. We can prove the following claim by relying on another pumping argument on corners:

Claim. *If $t_1 \cdots t_k$ is a path in K , then $t_1 t_2 \cdots t_k \mathbf{p}_{t_k} \cdots \mathbf{p}_{t_2} \mathbf{p}_{t_1}$ is a 0-cycle of $\mathcal{R}(\mathcal{G})$.*

Proof of Claim. We prove the property by induction on k . For $k = 1$, the property is immediate since $t_1 \mathbf{p}_{t_1}$ is a 0-cycle. Consider then k such that the property holds for k , and let us prove that it holds for $k+1$. We will exhibit two corner plays following $t_1 \cdots t_{k+1} \mathbf{p}_{t_{k+1}} \cdots \mathbf{p}_{t_1}$ of opposite weight and conclude with Lemma 4.1.

Let \mathbf{v}_0 be a corner of $\text{last}(t_{k+1})$. Since $t_{k+1} \mathbf{p}_{t_{k+1}}$ is a 0-cycle, there exists $w \in \mathbb{Z}$, a corner play ρ_0 following t_{k+1} ending in \mathbf{v}_0 with weight w and a corner play ρ'_0 following $\mathbf{p}_{t_{k+1}}$ beginning in \mathbf{v}_0 with weight $-w$. We name \mathbf{v}'_0 the corner of $\text{last}(t_k)$ where ρ'_0 ends. We consider any corner play ρ_1 following t_{k+1} from corner \mathbf{v}'_0 . The corner play $\rho'_0 \rho_1$ follows the path $\mathbf{p}_{t_{k+1}} t_{k+1}$ that is also a 0-cycle by Lemma 4.2, therefore ρ_1 has weight w . We denote by \mathbf{v}_1 the corner where ρ_1 ends. By iterating this construction, we obtain some corner plays $\rho_0, \rho_1, \rho_2, \dots$ following t_{k+1} and $\rho'_0, \rho'_1, \rho'_2, \dots$ following $\mathbf{p}_{t_{k+1}}$ such that ρ'_i goes from corner \mathbf{v}_i to \mathbf{v}'_i , and ρ_{i+1} from corner \mathbf{v}'_i to \mathbf{v}_{i+1} , for all $i \geq 0$. Moreover, all corner plays ρ_i have weight w and all corner plays ρ'_i have weight $-w$. Consider the first index ℓ such that $\mathbf{v}_\ell = \mathbf{v}_j$ for some $j < \ell$, which exists because the number of corners is finite.

We apply the induction hypothesis to find a corner play following $t_1 \cdots t_k \mathbf{p}_{t_k} \cdots \mathbf{p}_{t_1}$, going through the corner \mathbf{v}'_j in the middle: more formally, there exists w_α , a corner play ρ_α following $t_1 \cdots t_k$ ending in \mathbf{v}'_j with weight w_α and a corner play ρ'_α following $\mathbf{p}_{t_k} \cdots \mathbf{p}_{t_1}$ beginning in \mathbf{v}'_j with weight $-w_\alpha$. We apply the induction hypothesis a second time with corner $\mathbf{v}'_{\ell-1}$: there exists w_β , a corner play ρ_β following $t_1 \cdots t_k$ ending in $\mathbf{v}'_{\ell-1}$ with weight w_β and a corner play ρ'_β following $\mathbf{p}_{t_k} \cdots \mathbf{p}_{t_1}$ beginning in $\mathbf{v}'_{\ell-1}$ with weight $-w_\beta$.

The corner play $\rho_\alpha \rho_{j+1} \rho'_{j+1} \rho_{j+2} \rho'_{j+2} \cdots \rho'_{\ell-1} \rho'_\beta$, of weight $w_\alpha + (w - w)(\ell - j) - w_\beta = w_\alpha - w_\beta$, follows the cycle $t_1 \cdots t_k (t_{k+1} \mathbf{p}_{t_{k+1}})^{\ell-j} \mathbf{p}_{t_k} \cdots \mathbf{p}_{t_1}$. The corner play $\rho_\beta \rho_\ell \rho'_j \rho'_\alpha$, of weight $w_\beta + w - w - w_\alpha = w_\beta - w_\alpha$, follows the cycle $t_1 \cdots t_k t_{k+1} \mathbf{p}_{t_{k+1}} \mathbf{p}_{t_k} \cdots \mathbf{p}_{t_1}$. Since the game is almost-divergent, and those two corner plays are of opposite sign and in the same SCC, both have weight 0. The second corner play of weight 0 ensures that the cycle $t_1 \cdots t_{k+1} \mathbf{p}_{t_{k+1}} \cdots \mathbf{p}_{t_1}$ is a 0-cycle, by Lemma 4.1. \triangle

Now, if \mathbf{p} is a cycle of $\mathcal{R}(\mathcal{G})$ in K , there exists a cycle \mathbf{p}' such that $\mathbf{p}\mathbf{p}'$ is a 0-cycle. Since $\mathcal{R}(\mathcal{G})$ is an almost-divergent WTG, \mathbf{p} is a 0-cycle. \square

5.3. Values $-\infty$ coming from negative cycles. Equipped with the kernels, we are now ready to remove the only remaining configurations having a value $-\infty$ in $\mathcal{R}(\mathcal{G})$.

Proposition 5.5. *In an SCC of $\mathcal{R}(\mathcal{G})$, the set of configurations with value $-\infty$ is a union of regions computable in time linear in the size of $\mathcal{R}(\mathcal{G})$. Moreover, if we remove those states from $\mathcal{R}(\mathcal{G})$, the value of the other configurations does not change.*

Proof. If the SCC is non-negative, the cumulative weight cannot decrease along a cycle, thus, there can be no configurations with value $-\infty$.

If the SCC is non-positive, let T_t be the set of edges of $\mathcal{R}(\mathcal{G})$ whose end state has location in L_t . We prove that a configuration has value $-\infty$ if and only if it belongs to a region state where player Min can ensure the LTL formula on edges $\phi = (\mathsf{G} \neg T_t) \wedge \neg \mathsf{FG} T_K$: in simple

words, this means that Min can ensure to see an edge from T_t , or exit T_K if we enter the kernel at some point.

If (ℓ, r) is a region state where Min can ensure ϕ , Min can ensure a value of $-\infty$ from all configurations in (ℓ, r) by avoiding S_t for as long as desired, while not getting stuck in K , and thus going through an infinite number of negative cycles by Proposition 5.4. Conversely, suppose that (ℓ, r) is a region state from which Min cannot ensure ϕ . The winner of an ω -regular timed game only depends on the (finite) region abstraction [WTH92], and finite turn-based ω -regular games are determined [BL69]. Thus, we know that Max can ensure $\neg\phi = (FT_t) \vee FG T_K$. Then, from (ℓ, r) , Max must be able to reach S_t or stay in K forever. In both cases, Max can ensure a value above $-\infty$.

The procedure to detect $-\infty$ states thus consists in checking LTL formula $(G \neg T_t) \wedge \neg FG T_K$, and thus can be performed via three attractor computations in time linear in $|\mathcal{R}(\mathcal{G})|$. The set of region states with value $-\infty$ can then be removed safely from $\mathcal{R}(\mathcal{G})$, by Lemma 5.1 (since it is closed by attractor of Min). \square

From this point on, **we assume that no configurations of $\mathcal{R}(\mathcal{G})$ have value $+\infty$ or $-\infty$, and that the final weight function maps all configurations to \mathbb{R}** . Since fwt is piecewise affine with finitely many pieces, fwt is bounded. Let w_{\max}^t denote the supremum of $|\text{fwt}|$, ranging over all target configurations.

6. SEMI-UNFOLDING OF WTGS

Given an almost-divergent WTG \mathcal{G} , we describe the construction of its *semi-unfolding* $\mathcal{T}(\mathcal{G})$ (as depicted in Figure 11). This crucially relies on the absence of states with value $-\infty$ which has been explained in Section 5.

We only build the semi-unfolding $\mathcal{T}(\mathcal{G})$ of an SCC of \mathcal{G} starting from some state $(\ell_0, r_0) \in S$ of the region game, since it is then easy to glue all the semi-unfoldings together to get the one of the full game. **We thus assume in this section that $\mathcal{R}(\mathcal{G})$ is an SCC**. Since every configuration has finite value, we will prove that values of the game are bounded by $|\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t$. As a consequence, we can find a bound γ linear in $|\mathcal{R}(\mathcal{G})|$, w_{\max} and w_{\max}^t such that a play that visits some state outside the kernel more than γ times has weight strictly above $|\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t$, hence is useless for the value computation. This leads to considering the semi-unfolding $\mathcal{T}(\mathcal{G})$ of $\mathcal{R}(\mathcal{G})$ (nodes in the kernel are not unfolded, see Figure 11) such that each node not in the kernel is encountered at most γ times along a branch: the end of each branch is called a *stopped leaf* of the semi-unfolding. In particular, the depth of $\mathcal{T}(\mathcal{G})$ is bounded by $|\mathcal{R}(\mathcal{G})|\gamma$, and thus is polynomial in $|\mathcal{R}(\mathcal{G})|$, w_{\max} and w_{\max}^t . Leaves of the semi-unfolding are thus of two types: target leaves that are copies of target locations of $\mathcal{R}(\mathcal{G})$ for which we set the target weight as in $\mathcal{R}(\mathcal{G})$, and stop leaves for which we set their target weight as being constant to $+\infty$ if the SCC $\mathcal{R}(\mathcal{G})$ is non-negative, and $-\infty$ if the SCC is non-positive.

More formally, if (ℓ, r) is in K , we let $K_{\ell,r}$ be the part of K accessible from (ℓ, r) (note that $K_{\ell,r}$ is an SCC as K is a disjoint set of SCCs). We define the output edges of $K_{\ell,r}$ as being the output edges of K accessible from (ℓ, r) . If (ℓ, r) is not in K , the output edges of (ℓ, r) are the edges of $\mathcal{R}(\mathcal{G})$ starting in (ℓ, r) .

We define a tree T whose nodes are either labelled by region states $(\ell, r) \in S \setminus S_K$ or by kernels $K_{\ell,r}$, and whose edges are labelled by output edges in $\mathcal{R}(\mathcal{G})$. The root of the tree T is labelled with an initial region state (ℓ_0, r_0) , or K_{ℓ_0, r_0} (if (ℓ_0, r_0) belongs to the kernel),

and the successors of a node of T are then recursively defined by its output edges. When a state (ℓ, r) is reached by an output edge, the child is labelled by $K_{\ell, r}$ if $(\ell, r) \in \mathbf{K}$, otherwise it is labelled by (ℓ, r) . Edges in T are labelled by the edges used to create them. Along every branch, we stop the construction when either a final state is reached (*i.e.* a state not inside the current SCC) or the branch contains $3|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 2$ nodes labelled by the same state $((\ell, r)$ or $K_{\ell, r}$). Leaves of T with a location belonging to L_t are called *target leaves*, others are called *stopped leaves*.

We now transform T into a WTG $\mathcal{T}(\mathcal{G})$, by replacing every node labelled by a state (ℓ, r) by a different copy $(\tilde{\ell}, r)$ of (ℓ, r) . Those states are said to inherit from (ℓ, r) . Edges of T are replaced by the edges labelling them, and have a similar notion of inheritance. Every non-leaf node labelled by a kernel $K_{\ell, r}$ is replaced by a copy of the WTG $K_{\ell, r}$, output edges being plugged in the expected way. We deal with stopped leaves labelled by a kernel $K_{\ell, r}$ by replacing them with a single node copy of (ℓ, r) , like we dealt with node labelled by a state (ℓ, r) . State partition between players and weights are inherited from the copied states of $\mathcal{R}(\mathcal{G})$. The only initial state of $\mathcal{T}(\mathcal{G})$ is the state denoted by $(\tilde{\ell}_0, r_0)$ inherited from (ℓ_0, r_0) in the root of T (either (ℓ_0, r_0) or K_{ℓ_0, r_0}). The final states of $\mathcal{T}(\mathcal{G})$ are the states derived from leaves of T . If $\mathcal{R}(\mathcal{G})$ is a non-negative (respectively non-positive) SCC, the final weight function fwt is inherited from $\mathcal{R}(\mathcal{G})$ on target leaves and set to $+\infty$ (respectively $-\infty$) on stopped leaves.

The semi-unfolding of the WTG from Figure 10 can be found in Figure 18 of Appendix B.

Proposition 6.1. *Let \mathcal{G} be an almost-divergent WTG, and let (ℓ_0, r_0) be some region state of $\mathcal{R}(\mathcal{G})$. The semi-unfolding $\mathcal{T}(\mathcal{G})$ with initial state $(\tilde{\ell}_0, r_0)$ (a copy of a region state (ℓ_0, r_0)) is equivalent to \mathcal{G} , *i.e.* for all $\nu_0 \in r_0$, $\text{Val}_{\mathcal{G}}(\ell_0, \nu_0) = \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$.*

The rest of this section aims at proving Proposition 6.1, only in the case where $\mathcal{R}(\mathcal{G})$ is an SCC, since the general result then follows easily. First, we need information on the weight of finite plays in the region game.

Lemma 6.2. *All finite plays in $\mathcal{R}(\mathcal{G})$ have cumulative weight (ignoring final weights) at least $-|\mathcal{R}(\mathcal{G})|w_{\max}$ in the non-negative case, and at most $|\mathcal{R}(\mathcal{G})|w_{\max}$ in the non-positive case. Moreover, values of the game are bounded by $|\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t$.*

Proof. Suppose first that $\mathcal{R}(\mathcal{G})$ is a non-negative SCC. Consider a play ρ following a path \mathbf{p} . This path \mathbf{p} can be decomposed into $\mathbf{p} = \mathbf{p}_1 \mathbf{p}_1^c \cdots \mathbf{p}_k \mathbf{p}_k^c$ such that every \mathbf{p}_i^c is a cycle, and $\mathbf{p}_1 \cdots \mathbf{p}_k$ is a simple path in $\mathcal{R}(\mathcal{G})$ (thus $\sum_{i=1}^k |\mathbf{p}_i| \leq |\mathcal{R}(\mathcal{G})|$). Let us define all plays ρ_i and ρ_i^c as the restrictions of ρ on \mathbf{p}_i and \mathbf{p}_i^c . Now, since all plays following cycles have cumulative weight at least 0,

$$\text{wt}_{\Sigma}(\rho) = \sum_{i=1}^k \text{wt}_{\Sigma}(\rho_i) + \text{wt}_{\Sigma}(\rho_i^c) \geq \sum_{i=1}^k -w_{\max}|\rho_i| + 0 \geq -|\mathcal{R}(\mathcal{G})|w_{\max}$$

Similarly, we can show that every play in a non-positive SCC has cumulative weight at most $|\mathcal{R}(\mathcal{G})|w_{\max}$.

For the bound on the values, consider again two cases. If $\mathcal{R}(\mathcal{G})$ is a non-negative SCC, consider a positional attractor strategy σ_{Min} for Min toward S_t . Since all states have values below $+\infty$, all plays obtained from strategies of Max will follow simple paths of $\mathcal{R}(\mathcal{G})$, that have cumulative weight at most $|\mathcal{R}(\mathcal{G})|w_{\max}$ in absolute value. Similarly, if $\mathcal{R}(\mathcal{G})$ is a non-positive SCC, following the proof of Proposition 5.5, since all values are above $-\infty$,

Max can ensure $\neg\phi$, i.e. $(FT_t) \vee FGT_K$ on all states. Then we can construct a strategy σ_{Max} for Max combining an attractor strategy toward S_t on states satisfying FT_t , a safety strategy on states satisfying GT_K , and an attractor strategy toward the latter on all other states. Then, all plays obtained from strategies of Min will either not be winning (GT_K) or follow simple paths of $\mathcal{R}(\mathcal{G})$. Both cases imply that the values of the game are bounded by $|\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t$. \square

We can use this to obtain similar results in the semi-unfolding $\mathcal{T}(\mathcal{G})$.

Lemma 6.3. *All plays in $\mathcal{T}(\mathcal{G})$ from the initial state $(\tilde{\ell}_0, r_0)$ to a stopped leaf have cumulative weight at least $2|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 1$ if the SCC $\mathcal{R}(\mathcal{G})$ is non-negative, and at most $-2|\mathcal{R}(\mathcal{G})|w_{\max} - 2w_{\max}^t - 1$ if it is non-positive.*

Proof. Note that by construction, all finite paths in $\mathcal{T}(\mathcal{G})$ from the initial state to a stopped leaf can be decomposed as $\mathbf{p}'\mathbf{p}_1 \cdots \mathbf{p}_3|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 1$ with all \mathbf{p}_i being cycles around the same state. Additionally, those cycles cannot be 0-cycles by Proposition 5.4, since they take at least one edge outside of K . Therefore the restriction of ρ to $\mathbf{p}_1 \cdots \mathbf{p}_3|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 1$ has weight at least $3|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 1$ (in the non-negative case) and at most $-3|\mathcal{R}(\mathcal{G})|w_{\max} - 2w_{\max}^t - 1$ (in the non-positive case). The beginning of the play, following \mathbf{p}' , has cumulative weight at least $-|\mathcal{R}(\mathcal{G})|w_{\max}$ (in the non-negative case) and at most $|\mathcal{R}(\mathcal{G})|w_{\max}$ (in the non-positive case), by Lemma 6.2. \square

Consider two plays of the same length k in $\mathcal{R}(\mathcal{G})$ and $\mathcal{T}(\mathcal{G})$, respectively:

$$\begin{aligned} \rho &= ((\ell_1, r_1), \nu_1) \xrightarrow{d_1, t_1} \cdots \xrightarrow{d_{k-1}, t_{k-1}} ((\ell_k, r_k), \nu_k) \\ \tilde{\rho} &= ((\tilde{\ell}_1, r_1), \nu_1) \xrightarrow{d_1, \tilde{t}_1} \cdots \xrightarrow{d_{k-1}, \tilde{t}_{k-1}} ((\tilde{\ell}_k, r_k), \nu_k). \end{aligned}$$

They are said to *mimic* each other if every $(\tilde{\ell}_i, r_i)$ is inherited from (ℓ_i, r_i) and every edge \tilde{t}_i is inherited from the edge t_i of $\mathcal{R}(\mathcal{G})$. Combining Lemmas 6.2 and 6.3, we obtain:

Lemma 6.4. *If $\mathcal{R}(\mathcal{G})$ is a non-negative (respectively non-positive) SCC, every play from the initial state and with cumulative weight less than $|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 1$ (respectively greater than $-|\mathcal{R}(\mathcal{G})|w_{\max} - 2w_{\max}^t - 1$) can be mimicked in $\mathcal{T}(\mathcal{G})$ without reaching a stopped leaf. Conversely, every play in $\mathcal{T}(\mathcal{G})$ reaching a target leaf can be mimicked in $\mathcal{R}(\mathcal{G})$.*

Proof. We prove only the non-negative case, since the other case is symmetrical. Let ρ be a play of $\mathcal{R}(\mathcal{G})$ with cumulative weight less than $|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 1$. Consider the branch of the unfolded game it follows. If ρ cannot be mimicked in $\mathcal{T}(\mathcal{G})$, then a prefix of ρ reaches the stopped leaf of that branch when mimicked in $\mathcal{T}(\mathcal{G})$. In this situation, ρ starts by a prefix of weight at least $2|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 1$ by Lemma 6.3 and then ends with a suffix play of weight at least $-|\mathcal{R}(\mathcal{G})|w_{\max}$ by Lemma 6.2, and that contradicts the initial assumption. The converse is true by construction. \square

Then, intuitively, the plays of $\mathcal{R}(\mathcal{G})$ starting in an initial configuration that cannot be mimicked in $\mathcal{T}(\mathcal{G})$ are not useful for value computation. To obtain Proposition 6.1, we now prove that for all valuations $\nu_0 \in r_0$, $\text{Val}_{\mathcal{G}}(\ell_0, \nu_0) = \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$. By Lemma 2.13, we already know that $\text{Val}_{\mathcal{G}}(\ell_0, \nu_0) = \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell_0, r_0), \nu_0)$. Recall that we only left finite values in $\mathcal{R}(\mathcal{G})$ (in the final weight functions, in particular), and more precisely $|\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell_0, r_0), \nu_0)| \leq |\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t$ by Lemma 6.2. We first show that the value is also finite in $\mathcal{T}(\mathcal{G})$. Indeed, if $\text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) = +\infty$, since we assumed all final weights of $\mathcal{R}(\mathcal{G})$ bounded, we are necessarily in the non-negative case, and Max is able to ensure stopped leaves reachability.

Claim. If $\text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) = +\infty$, then there are no strategies in $\mathcal{R}(\mathcal{G})$ for Min ensuring weight less than $|\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t + 1$ from (ℓ_0, r_0) .

Thus, we can obtain the contradiction $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell_0, r_0), \nu_0) > |\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t$.

Proof of Claim. By contradiction, consider a strategy σ_{Min} of Min ensuring weight $A \leq |\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t + 1$ in $\mathcal{R}(\mathcal{G})$. Then, for all σ_{Max} , the cumulative weight of the play $\text{play}_{\mathcal{R}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0, \sigma_{\text{Min}}, \sigma_{\text{Max}})$ (reaching target configuration (ℓ, ν)) is at most equal to $A - \text{fwt}(\ell, \nu) \leq |\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 1$, and by Lemma 6.4 this play does not reach a stopped leaf when mimicked in $\mathcal{T}(\mathcal{G})$, which is absurd. \triangle

If $\text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) = -\infty$, we are necessarily in the non-positive case, and, by construction, this implies that Min ensures stopped leaves reachability in $\mathcal{T}(\mathcal{G})$.

Claim. If $\text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) = -\infty$, then there are no strategies in $\mathcal{R}(\mathcal{G})$ for Max ensuring weight above $-|\mathcal{R}(\mathcal{G})|w_{\max} - w_{\max}^t - 1$ from (ℓ_0, r_0) .

Thus, we can obtain the contradiction $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell_0, r_0), \nu_0) < -|\mathcal{R}(\mathcal{G})|w_{\max} - w_{\max}^t$.

Proof of Claim. By contradiction, consider a strategy σ_{Max} of Max ensuring weight $A \geq -|\mathcal{R}(\mathcal{G})|w_{\max} - w_{\max}^t - 1$ in $\mathcal{R}(\mathcal{G})$. Then, for all σ_{Min} , the cumulative weight of the play $\text{play}_{\mathcal{R}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0, \sigma_{\text{Min}}, \sigma_{\text{Max}})$ (reaching target configuration (ℓ, ν)) is at least $A - \text{fwt}(\ell, \nu)$, thus at least $-|\mathcal{R}(\mathcal{G})|w_{\max} - 2w_{\max}^t - 1$. By Lemma 6.4, this play does not reach a stopped leaf when mimicked in $\mathcal{T}(\mathcal{G})$, which is absurd. \triangle

If $\mathcal{R}(\mathcal{G})$ is non-negative, for all $\varepsilon > 0$ we can fix an ε -optimal strategy σ_{Min} for Min in $\mathcal{T}(\mathcal{G})$. Every play derived from σ_{Min} in $\mathcal{T}(\mathcal{G})$ reaches a target leaf, and can thus be mimicked in $\mathcal{R}(\mathcal{G})$ by Lemma 6.4. Therefore, σ_{Min} can be mimicked in $\mathcal{R}(\mathcal{G})$, where it keeps the same value. From this we deduce $\text{Val}_{\mathcal{R}(\mathcal{G})}((\ell_0, r_0), \nu_0) \leq \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$. If $\mathcal{R}(\mathcal{G})$ is non-positive, the same reasoning applies by considering an ε -optimal strategy for Max in $\mathcal{T}(\mathcal{G})$.

Let us now show the reverse inequality. If $\mathcal{R}(\mathcal{G})$ is non-negative, let us fix $0 < \varepsilon < 1$, an ε -optimal strategy σ_{Min} for Min in $\mathcal{R}(\mathcal{G})$, and a strategy σ_{Max} of Max in $\mathcal{R}(\mathcal{G})$. Let ρ be their outcome $\text{play}_{\mathcal{R}(\mathcal{G})}((\ell_0, r_0), \nu_0, \sigma_{\text{Min}}, \sigma_{\text{Max}})$, ρ_k be the finite prefix of ρ defining its cumulative weight and (ℓ_k, ν_k) be the configuration defining its final weight, such that $\text{wt}_{\mathcal{R}(\mathcal{G})}(\rho) = \text{wt}_{\Sigma}(\rho_k) + \text{fwt}(\ell_k, \nu_k)$. Then,

$$\text{wt}_{\mathcal{R}(\mathcal{G})}(\rho) \leq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell_0, r_0), \nu_0) + \varepsilon < |\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t + 1$$

therefore

$$\text{wt}_{\Sigma}(\rho_k) < |\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t + 1 - \text{fwt}(\ell_k, \nu_k) \leq |\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 1$$

and by Lemma 6.4 all such plays ρ can be mimicked in $\mathcal{T}(\mathcal{G})$, so that

$$\text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) \leq \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell_0, r_0), \nu_0)$$

Once again, if $\mathcal{R}(\mathcal{G})$ is non-positive, the same reasoning applies by considering an ε -optimal strategy for Max in $\mathcal{R}(\mathcal{G})$. This ends the proof of Proposition 6.1.

Remark 6.5. Note that the semi-unfolding procedure of an SCC depends on w_{\max}^t , where fwt can be the value function of an SCCs under the current one. Assuming all configurations have finite values, we can extend the reasoning of Lemma 6.2 and bound all values in the

full game by $|\mathcal{R}(\mathcal{G})|w_{\max} + w_{\max}^t$, which lets us bound uniformly the unfolding depth of each SCC and gives us a bound on the depth of the complete semi-unfolding tree:

$$|\mathcal{R}(\mathcal{G})|(3|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 2) + 1 \quad (6.1)$$

7. COMPUTING VALUES FOR ACYCLIC WTGS

In this section, we focus on the class of acyclic WTGs where the value problem is decidable, as shown by [ABM04]. In our context, this will give us a way to compute the value of the semi-unfolding $\mathcal{T}(\mathcal{G})$ of an almost-divergent WTG that contains no kernels: this is equivalent for the WTG to be divergent indeed. This section is particularly technical, and can thus be skipped by non-interested readers. We have chosen to give such detailed explanations of the techniques used in [ABM04] (with independent proofs) for several reasons. On the one hand, our setting is more general, in the sense that we allow for negative weights and for final weights, where the authors of [ABM04] do not do so explicitly. On the other hand, their result is stated for concurrent games, a generalisation of the turn-based games we consider. This leads to simplifications in the proofs, and makes easier some parts of the complexity analysis. We will need, in Section 8, to bound the partial derivatives of the functions we compute. This cannot be deduced from their result directly. We present their techniques in a new, more symbolic light, by performing computations on the entire state-space at once instead of region by region. For reasons detailed in Section 7.4 and in [BG19], we are not able to replicate their (incomplete) complexity analysis. We will therefore rely on a doubly-exponential upper bound instead of the exponential one claimed in [ABM04]. Last but not least, this detailed analysis allows us to solve the synthesis of ε -optimal strategies for both players, as will be detailed in Section 10.

The main result of this section is a symbolic algorithm for computing the value $\mathbf{V}^i = \text{Val}_{\mathcal{G}}^i$ defined in Section 2.8.

Theorem 7.1. *Given $i \geq 0$, computing $\text{Val}_{\mathcal{G}}^i$ can be done in time doubly-exponential in i and exponential in the size of \mathcal{G} .*

The intuition behind the result is the observation that the mappings \mathbf{V}_{ℓ}^0 are piecewise affine for all ℓ , and a proof that \mathcal{F} preserves piecewise affinity, so that all iterates \mathbf{V}_{ℓ}^i can be computed using piecewise affine functions. In order to bound the size of \mathbf{V}_{ℓ}^i (in particular, its number of pieces), we need the fine encoding via cells and partition functions defined in Section 2.3.

7.1. About complexity bounds. In this section we will assume without loss of generality that the number of clocks $n \geq 1$. The case $n = 0$ (finite weighted games) will be detailed in Section 11. The piecewise affine value function \mathbf{V} is encoded as pairs $(P_{\ell}, F_{\ell})_{\ell \in L}$ such that $\llbracket F_{\ell} \rrbracket = \mathbf{V}_{\ell}$ for all $\ell \in L$. We detail now how we measure the complexity of a pair (P, F) . As mentioned before, we assume without loss of generality that (in)equalities involved in the definition of cells only use integers. This is not the case for value functions which are described by equations $y = a_1x_1 + \dots + a_nx_n + b$ with a_i and b in \mathbb{Q} . In order to track their size, we instead write $a_yy = a_1x_1 + \dots + a_nx_n + b$, with all a_i and b integers of \mathbb{Z} , and $a_y \in \mathbb{N}_{>0}$.

Definition 7.2. Let (P, F) be a pair composed of a partition P and a value function F defined on that partition. The *complexity* of (P, F) is the pair $\langle m, \beta \rangle$ where the *size complexity* m and the *constant complexity* β are defined as follows. First, we define the size complexity as $m = |\mathcal{E}_P| + m_0$, where m_0 denotes the number of inequalities in the encoding of the domain c_P .¹² Second, the constant complexity β is the smallest natural number that upper bounds (in absolute value) every finite constant (partial derivatives and additive constants) in the affine (in)equalities of P and F . Given a piecewise affine value function (P, F) of complexity $\langle m, \beta \rangle$, we say that (P, F) has complexity at most $\langle m', \beta' \rangle$ if $m \leq m'$ and $\beta \leq \beta'$.

From this definition, we can bound the memory needed to store a partition function:

Lemma 7.3. *Let \mathbf{V} be represented by a pair (P, F) of complexity at most $\langle m, \beta \rangle$. Then it can be represented using a space in $\mathcal{O}(2^n(n+2)(m+1)^{n+1} \lceil \log \beta \rceil)$.*

Proof. Storing P requires storing m affine expressions, each stored in space $(n+1) \lceil \log \beta \rceil$. Moreover, each affine expression of the value function F , one for each base cell of P , is stored in space $(n+2) \lceil \log \beta \rceil$. P has at most $2^n(m+1)^n$ base cells by Lemma 2.6, and storing each base cell explicitly requires storing at most m borders per base cell. In total, \mathbf{V} can be represented in space $2^{n+1}(n+2)(m+1)^{n+1} \lceil \log \beta \rceil$. \square

7.2. Operations over value functions. Our goal is now to compute the value function $\mathcal{F}(\mathbf{V})$ encoded by $(P'_\ell, F'_\ell)_{\ell \in L}$, when \mathbf{V} is encoded by $(P_\ell, F_\ell)_{\ell \in L}$. We decompose the computation into smaller operations that we first introduce and compute over partitions, while explaining how they affect the complexity parameter $\langle m, \beta \rangle$.

If P_1 and P_2 are partitions over the same domain c_P , let $P_1 \oplus P_2$ denote the coarsest partition consistent with both P_1 and P_2 : each base cell c_b of $P_1 \oplus P_2$ corresponds to an intersection $c_1 \cap c_2$, with c_1 a base cell of P_1 and c_2 a base cell of P_2 . It is obtained from $\mathcal{E}_{P_1 \oplus P_2} = \mathcal{E}_{P_1} \cup \mathcal{E}_{P_2}$. Note that if $P_1 \dots P_q$ are partitions of complexity at most $\langle m, \beta \rangle$, $P_1 \oplus \dots \oplus P_q$ is a partition of size complexity at most $\langle qm, \beta \rangle$. Now, the minimum (respectively maximum) of a finite set of piecewise affine value functions can be computed with partitions.

Lemma 7.4 [ABM04, Thm. 1]. *Let $(P_i, F_i)_{1 \leq i \leq q}$ be piecewise affine value functions, defined over the same domain c_P , where each (P_i, F_i) has complexity at most $\langle m, \beta \rangle$. There exists a piecewise affine value function (P', F') of domain c_P and complexity at most $\langle m', \beta' \rangle$, where $m' = qm + q^2$, and $\beta' = 2\beta^2$, such that $\llbracket F' \rrbracket = \min_{1 \leq i \leq q} \llbracket F_i \rrbracket$ (respectively $\llbracket F' \rrbracket = \max_{1 \leq i \leq q} \llbracket F_i \rrbracket$).*

Proof. Let P' be $P_1 \oplus \dots \oplus P_q$. Let c denote a base cell in P' , corresponding to an intersection $c_1 \cap \dots \cap c_q$ of base cells of P_1, \dots, P_q respectively. Consider the affine value functions $F_1(c_1), \dots, F_q(c_q)$. Each of these is defined by an equation of the form $a_y y = a_1 x_1 + \dots + a_n x_n + b$, that can be seen as affine equalities over variables $\mathcal{X} \uplus \{y\}$, denoted E_1, \dots, E_q , or equivalently as sets of points in $\mathbb{R}^{\mathcal{X} \uplus \{y\}}$, denoted $\llbracket E_1 \rrbracket, \dots, \llbracket E_q \rrbracket$. If E and E' are such equalities, of equations $a_y y = a_1 x_1 + \dots + a_n x_n + b$ and $a'_y y = a'_1 x_1 + \dots + a'_n x_n + b'$, the intersection $\llbracket E \rrbracket \cap \llbracket E' \rrbracket$ is either empty or, by elimination of y , it satisfies the equation

$$(a_y a'_1 - a'_y a_1) x_1 + \dots + (a_y a'_n - a'_y a_n) x_n + (a_y b' - a'_y b) = 0.$$

¹²If $c_P = [0, M]^x$, then $m_0 = n$.

This describes an affine equality over \mathcal{X} , that we denote $E \cap_y E'$, with constant complexity at most $2\beta^2$. Now, let us partition $P_1 \oplus \dots \oplus P_q$ by the set of all such intersections

$$\{E_i \cap_y E_j \mid 1 \leq i, j \leq q \wedge \llbracket E_i \rrbracket \cap \llbracket E_j \rrbracket \neq \emptyset\}.$$

On every sub-cell c' of this partition, there exists $j \in \{1, \dots, n\}$ such that for every $\nu \in c'$, $\llbracket F_j \rrbracket(\nu) = \min_{1 \leq i \leq q} \llbracket F_i \rrbracket(\nu)$. Therefore, we define F' on c' as equal to $F_j(c_j)$. The partition $P_1 \oplus \dots \oplus P_q$ has size complexity at most qm , and we added at most q^2 intersections $E \cap_y E'$, resulting in a partition of the expected complexity. \square

For all $\nu \in [0, M]^{\mathcal{X}}$, let $d_\nu = \sup\{d \mid \nu + d \in [0, M]^{\mathcal{X}}\} \in \mathbb{R}_{\geq 0}$ denote the greatest valid delay from ν (in fact, $d_\nu = M - \|\nu\|_\infty$). Consider the following operations:

- If $\mathcal{Y} \subseteq \mathcal{X}$ is a set of clocks and $\ell \in L$, let $\text{Unreset}_{\mathcal{Y}}(\mathbf{V}_\ell) : [0, M]^{\mathcal{X}} \rightarrow \mathbb{R}_\infty$ denote the value function such that for all ν ,

$$\text{Unreset}_{\mathcal{Y}}(\mathbf{V}_\ell)(\nu) = \mathbf{V}_\ell(\nu[\mathcal{Y} := 0]).$$

- If g is a guard over \mathcal{X} and $\ell \in L$, let $\text{Guard}_g(\mathbf{V}_\ell) : [0, M]^{\mathcal{X}} \rightarrow \mathbb{R}_\infty$ denote the value function such that for all ν ,

$$\text{Guard}_g(\mathbf{V}_\ell)(\nu) = \begin{cases} \mathbf{V}_\ell(\nu) & \text{if } \nu \models g \\ -\infty & \text{if } \nu \not\models g \wedge \ell \in L_{\text{Max}} \\ +\infty & \text{if } \nu \not\models g \wedge \ell \in L_{\text{Min}}. \end{cases}$$

The values $+\infty$ and $-\infty$ in Guards_g ensure that players cannot choose invalid delays: By the no-deadlocks assumption, from every configuration there exists a transition in $\llbracket \mathcal{G} \rrbracket$, whose value will win against $+\infty$ or $-\infty$ in the subsequent equality (7.2).

- If $e \in E$ is an edge from ℓ to ℓ' , let $\text{Pre}_e(\mathbf{V}_{\ell'}) : [0, M]^{\mathcal{X}} \rightarrow \mathbb{R}_\infty$ denote the value function such that for all ν ,

$$\text{Pre}_e(\mathbf{V}_{\ell'}) (\nu) = \begin{cases} \sup_{d \in [0, d_\nu]} [d \cdot \text{wt}(\ell) + \text{wt}(e) + \mathbf{V}_{\ell'}(\nu + d)] & \text{if } \ell \in L_{\text{Max}} \\ \inf_{d \in [0, d_\nu]} [d \cdot \text{wt}(\ell) + \text{wt}(e) + \mathbf{V}_{\ell'}(\nu + d)] & \text{if } \ell \in L_{\text{Min}}. \end{cases} \quad (7.1)$$

Then, if $\mathbf{V}' = \mathcal{F}(\mathbf{V})$, it holds that

$$\mathbf{V}'_\ell = \begin{cases} \mathbf{V}_\ell & \text{if } \ell \in L_t \\ \max_{e=(\ell, g, \mathcal{Y}, \ell')} \text{Pre}_e(\text{Guard}_g(\text{Unreset}_{\mathcal{Y}}(\mathbf{V}_{\ell'}))) & \text{if } \ell \in L_{\text{Max}} \\ \min_{e=(\ell, g, \mathcal{Y}, \ell')} \text{Pre}_e(\text{Guard}_g(\text{Unreset}_{\mathcal{Y}}(\mathbf{V}_{\ell'}))) & \text{if } \ell \in L_{\text{Min}} \setminus L_t \end{cases} \quad (7.2)$$

where e ranges over the edges in \mathcal{G} that start from ℓ . We have detailed in Lemma 7.4 how one can perform the min and max operations over partitions. Let us now focus on the Guard_g and $\text{Unreset}_{\mathcal{Y}}$ operations.

Lemma 7.5. *Let (P, F) be a piecewise affine value function of complexity at most $\langle m, \beta \rangle$. Let g be a non-diagonal guard in \mathcal{G} . Then there exists a piecewise affine value function (P', F') of complexity at most $\langle m', \beta' \rangle$, where $m' = m + 2n$ and $\beta' = \max(\beta, M)$, such that $\llbracket F' \rrbracket = \text{Guard}_g(\llbracket F \rrbracket)$.*

Proof. The non-diagonal guard g can be encoded as a cell $I_1 \wedge \dots \wedge I_{2n}$, with one upper and one lower inequality for each clock. We define P' from P with the set of affine equalities $\mathcal{E}_P \cup \{E(I_1), \dots, E(I_{2n})\}$. It follows that each base cell of P' is either entirely included in g or entirely outside of it. We can thus define F' appropriately, such that $\llbracket F' \rrbracket = \text{Guard}_g(\llbracket F \rrbracket)$. \square

We continue our study with the reset of clocks, that we illustrate in Figure 15.

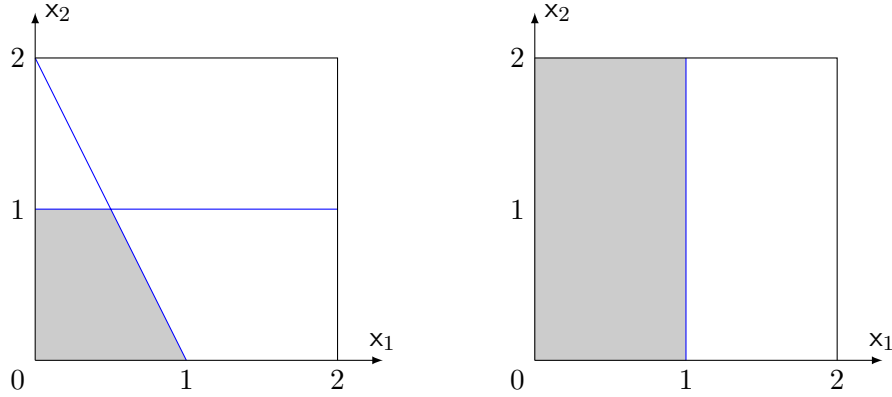


Figure 15: On the left, the partition of Figure 5. On the right, the corresponding partition obtained by applying Lemma 7.6 for reset $\mathcal{Y} = \{x_2\}$. The affine value function of the grey cell on the right is obtained from the grey cell on the left, by setting the partial derivative of x_2 to 0.

Lemma 7.6. *Let (P, F) be a piecewise affine value function of complexity at most $\langle m, \beta \rangle$. Let \mathcal{Y} be a set of clocks. Then there exists a piecewise affine value function (P', F') of complexity at most $\langle m, \beta \rangle$ such that $\llbracket F' \rrbracket = \text{Unreset}_{\mathcal{Y}}(\llbracket F \rrbracket)$.*

Proof. If $E : a_1x_1 + \dots + a_nx_n + b = 0$ is an affine equality, let $\text{Unreset}_{\mathcal{Y}}(E)$ denote the affine equality $a'_1x_1 + \dots + a'_nx_n + b = 0$, with $a'_i = 0$ if $x_i \in \mathcal{Y}$ and $a'_i = a_i$ otherwise, for $i \in [1, n]$. We extend this operator to affine inequalities I in the same way. For each valuation $\nu \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$, and E an affine equality (respectively inequality) $\nu \models \text{Unreset}_{\mathcal{Y}}(E)$ if and only if $\nu[\mathcal{Y} := 0] \models E$. Then, if $c = I_1 \wedge \dots \wedge I_p$ is a cell, let $\text{Unreset}_{\mathcal{Y}}(c)$ denote the cell $\text{Unreset}_{\mathcal{Y}}(I_1) \wedge \dots \wedge \text{Unreset}_{\mathcal{Y}}(I_p)$. It follows that $\nu \in \text{Unreset}_{\mathcal{Y}}(c)$ if and only if $\nu[\mathcal{Y} := 0] \in c$. In particular, if c does not intersect the sub-space where every clock in \mathcal{Y} equals 0, then $\text{Unreset}_{\mathcal{Y}}(c) = \emptyset$.

Similarly, if c is a base cell of P and F maps c to the affine value function $y = a_1x_1 + \dots + a_nx_n + b$, let $\text{Unreset}_{\mathcal{Y}}(F(c))$ denote the affine function $y = a'_1x_1 + \dots + a'_nx_n + b$ with for $i \in [1, n]$, $a'_i = 0$ if $x_i \in \mathcal{Y}$ and $a'_i = a_i$ otherwise. Then, for every $\nu \in \text{Unreset}_{\mathcal{Y}}(c)$, it holds that $\mathbf{V}(\nu[\mathcal{Y} := 0]) = \text{Unreset}_{\mathcal{Y}}(F(c))(\nu)$.

We construct P' from P by replacing c_P by $\text{Unreset}_{\mathcal{Y}}(c_P)$, and by replacing $\mathcal{E}_P = \{E_1, E_2, \dots\}$ by $\{\text{Unreset}_{\mathcal{Y}}(E_1), \text{Unreset}_{\mathcal{Y}}(E_2), \dots\}$. If c is a base cell of P , and $\text{Unreset}_{\mathcal{Y}}(c)$ is non-empty, we let $F'(c) = \text{Unreset}_{\mathcal{Y}}(F(c))$. The result is a partition P' with the desired complexity, and a partition function F' such that $\llbracket F' \rrbracket = \text{Unreset}_{\mathcal{Y}}(\llbracket F \rrbracket)$. \square

7.3. Tubes and diagonals. All that is left is the Pre_e operation. It is more challenging, and requires some extra machinery, that we now detail, related to *diagonal* behaviours that naturally arise when dealing with time-elapses.

An affine inequality (respectively equality) is *diagonal* if the sum of its partial derivatives is null, i.e. $a_1 + \dots + a_n = 0$. It follows that if ν satisfies a diagonal I then $\nu + d \models I$ for all $d \in \mathbb{R}$. A cell is called a *tube* when all of its inequalities are diagonal. When the cell is a sub-cell of some domain c_P in a partition P , we relax this definition slightly, to allow for non-diagonal borders inherited from c_P .

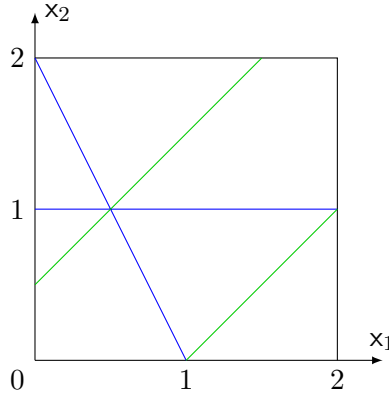


Figure 16: The atomic tube partition derived from the partition of Figure 5 by Lemma 7.7.

A *tube partition* is a partition P where the set \mathcal{E}_P of affine equalities is split between the set \mathcal{E}_P^d of diagonal affine equalities and the set \mathcal{E}_P^{nd} of non-diagonal ones. A tube partition induces a partition of c_P by \mathcal{E}_P^d into base cells that are tubes, called the *base tubes* of P .

Given two affine equalities $E : a_1x_1 + \dots + a_nx_n + b = 0$ and $E' : a'_1x_1 + \dots + a'_nx_n + b' = 0$, let $A = a_1 + \dots + a_n$ and $A' = a'_1 + \dots + a'_n$ denote the sums of their respective partial derivatives. We define their diagonal intersection as

$$E \cap_d E' : (Aa'_1 - A'a_1)x_1 + \dots + (Aa'_n - A'a_n)x_n + (Ab' - A'b) = 0$$

Observe that $E \cap_d E'$ is a (possibly empty) diagonal equality that contains $E \cap E'$. Moreover, if E (respectively E') is diagonal then $E \cap_d E'$ is equivalent to E (respectively E'). Now, given a cell $c = I_1 \wedge \dots \wedge I_m$ and a set \mathcal{E} of affine equalities, let $\bar{\mathcal{E}}$ denote¹³ $\mathcal{E} \cup \{E(I_1), \dots, E(I_m)\}$, and let $\text{Tube}(c, \mathcal{E})$ denote $\{E \cap_d E' \mid E, E' \in \bar{\mathcal{E}}\}$. The pair (c, \mathcal{E}) is said *atomic* if c is partitioned by $\text{Tube}(c, \mathcal{E})$ in only one cell (equal to c). Intuitively, (c, \mathcal{E}) is atomic if the affine equalities in \mathcal{E} and in the borders of c do not intersect within the smallest tube that contains c .

A tube partition is *atomic* if for every base tube c in the partition of c_P by \mathcal{E}_P^d , the pair (c, \mathcal{E}_P^{nd}) is atomic. Intuitively, this means that in the non-diagonal part of P , the equalities that split cells into sub-cells do not intersect within their base tube. Tube partitions can be made atomic, by introducing a bounded amount of diagonal affine equalities.

Lemma 7.7 [ABM04, Lem. 3]. *Let (P, F) be a piecewise affine value function of complexity at most $\langle m, \beta \rangle$. Then there exists a piecewise affine value function (P', F') where P' is an atomic tube partition, of complexity at most $\langle m', \beta' \rangle$, where¹⁴ $m' = m + m^2 \text{Splits}(m, n)$ and $\beta' = 2n\beta^2$, such that $\llbracket F' \rrbracket = \llbracket F \rrbracket$.*

Proof. We add a set of new diagonal equalities, that contain all equalities $E \cap_d E'$ derived from the base cells included in base tubes of P . Then, there are at most m^2 new diagonals for each base cell, and the resulting tube partition must be atomic, as any $E \cap_d E'$ belongs to $\bar{\mathcal{E}}$, and base cells of P cannot be partitioned by $\bar{\mathcal{E}}$. As $|A|, |A'| \leq n\beta$, we can bound by $2n\beta^2$ the constants in inequalities $E \cap_d E'$. \square

¹³Given an affine inequality I , $E(I)$ denotes the associated affine equality, see page 10.

¹⁴ $\text{Splits}(m, n)$ has been defined on page 10.

Figure 16 represents the atomic tube partition P' associated to the partition P displayed in Figure 5. The tube partition P' has 2 diagonal equalities in green, resulting in 5 base tubes. The result P' is therefore an atomic tube partition of complexity $\langle 6, 2 \rangle$.¹⁵

We can now compute $\text{Pre}_e(\mathbf{V}_{\ell'})$, with $\mathbf{V}_{\ell'}$ a value function encoded as a tube partition (P, F) , and $e \in E$ an edge from ℓ to ℓ' , using (7.1). We will assume in the following that ℓ is a location of **Max**, but the case of **Min** is symmetrical. Let us fix a valuation $\nu \in [0, M)^{\mathcal{X}}$. For every delay $d \in [0, d_\nu]$, consider the term $\mathbf{V}_{\ell'}(\nu + d)$ of (7.1). The valuations $\nu + d$ belong to a diagonal line of $\mathbb{R}_{\geq 0}^{\mathcal{X}}$, and range from ν to $\nu + d_\nu$. The segment $[\nu, \nu + d_\nu]$ intersects a finite set \mathcal{C}_ν of base cells in P . For each such cell c , we isolate from the segment $[\nu, \nu + d_\nu]$ two delays:

$$d_1 = \inf\{d \in [0, d_\nu] \mid \nu + d \in c\} \quad \text{and} \quad d_2 = \sup\{d \in [0, d_\nu] \mid \nu + d \in c\}$$

As $d \mapsto \llbracket F \rrbracket(\nu + d)$ is affine over c , so is $d \mapsto d \cdot \text{wt}(\ell) + \text{wt}(e) + \llbracket F \rrbracket(\nu + d)$, thus the supremum of $d \cdot \text{wt}(\ell) + \text{wt}(e) + \llbracket F \rrbracket(\nu + d)$ for $\nu + d \in c$ must either be reached at (or arbitrarily close to) d_1 , or at (or arbitrarily close to) d_2 . Note that $\nu + d_2$ must belong to a non-diagonal border of c , while $\nu + d_1$ either belongs to a non-diagonal border of c or equals ν (whenever $\nu \in c$). Thus, the optimal value of d for evaluating the supremum must correspond to either delay 0 or to a delay leading ν to a non-diagonal border (this is also proven in [ABM04]).

If B is a non-diagonal border of c , and ν is a valuation of $\mathbb{R}_{\geq 0}^{\mathcal{X}}$, there exists a unique $d \in \mathbb{R}$ such that $\nu + d \in \llbracket B \rrbracket$. In fact, if B is described by $a_1x_1 + \dots + a_nx_n + b = 0$ and $A = a_1 + \dots + a_n$, then

$$d = -\frac{a_1 \cdot \nu(x_1) + \dots + a_n \cdot \nu(x_n) + b}{A}$$

We denote this delay $d_{\nu, B}$. Observe that it must belong to $[0, d_\nu]$ as $\llbracket B \rrbracket$ is reachable from ν by time-elapse.¹⁶

If c is a cell of \mathcal{C}_ν , let $\mathcal{B}_\nu(c)$ denote the non-diagonal borders of c reachable from ν by time-elapse. The supremum $\text{Pre}_e(\llbracket F \rrbracket)(\nu)$ is then equal to

$$\max \left(\llbracket F \rrbracket(\nu), \max_{c \in \mathcal{C}_\nu} \max_{B \in \mathcal{B}_\nu(c)} [d_{\nu, B} \cdot \text{wt}(\ell) + \text{wt}(e) + \llbracket F(c) \rrbracket(\nu + d_{\nu, B})] \right)$$

where $\llbracket F \rrbracket(\nu)$ corresponds to the delay 0, and $\llbracket F(c) \rrbracket(\nu + d_{\nu, B})$ corresponds to a jump in cell c arbitrarily close to B .¹⁷

If the tube partition (P, F) is atomic, it follows that every other valuation in the same cell $\nu' \in [\nu]_P$ can reach the same set of borders by time elapse, *i.e.* $\mathcal{C}_{\nu'} = \mathcal{C}_\nu$ and $\mathcal{B}_{\nu'}(c) = \mathcal{B}_\nu(c)$ for all $c \in \mathcal{C}_\nu$. As a result, we rename those sets $\mathcal{C}_{c'}$ and $\mathcal{B}_{c'}(c)$ if $c' = [\nu]_P$ and $c \in \mathcal{C}_{c'}$. We introduce an operator $\text{Pre}_{e, c, B}$, indexed by an edge, a cell and a non-diagonal border of the cell, that maps a partition function F to the value function $\nu \mapsto d_{\nu, B} \cdot \text{wt}(\ell) + \text{wt}(e) + \llbracket F(c) \rrbracket(\nu + d_{\nu, B})$. As a consequence, for each base cell c_b of P , $\text{Pre}_e(\llbracket F \rrbracket)$ restricted to domain c_b equals

$$\max(\llbracket F \rrbracket, \max_{c \in \mathcal{C}_{c_b}} \max_{B \in \mathcal{B}_{c_b}(c)} \text{Pre}_{e, c, B}(F)) \tag{7.3}$$

Lemma 7.8. *Let (P, F) be a piecewise affine value function of complexity at most $\langle m, \beta \rangle$, where P is an atomic tube partition. Let c_b be a base cell of P , e be an edge from ℓ to ℓ' ,*

¹⁵the size complexity includes the two borders of the domain

¹⁶Note that it can be equal to d_ν , as $x - M = 0$ is a border of the cell $[0, M)^{\mathcal{X}}$.

¹⁷In particular, if $\nu + d_{\nu, B} \notin c$ then $F(c)$ evaluated on $\nu + d_{\nu, B}$ may not equal $\llbracket F \rrbracket(\nu + d_{\nu, B})$.

c be a cell in \mathcal{C}_{c_b} and B be a border in $\mathcal{B}_{c_b}(c)$. Then there exists an affine value function f with constants bounded by $\beta' = 4n\beta^2 \max(w_{\max}^L, w_{\max}^E)$, such that $f = \text{Pre}_{e,c,B}(F)$ on c_b .

Proof. Let ν be a valuation in c_b . As $B \in \mathcal{B}_{c_b}(c)$ it holds that $d_{\nu,B} \in [0, d_\nu]$, such that

$$\text{Pre}_{e,c,B}(F)(\nu) = d_{\nu,B} \cdot \text{wt}(\ell) + \text{wt}(e) + \llbracket F(c) \rrbracket(\nu + d_{\nu,B}).$$

Let $a_y y = a_1 x_1 + \dots + a_n x_n + b$ be the equation of $F(c)$, and let $A = a_1 + \dots + a_n$. Let $a'_1 x_1 + \dots + a'_n x_n + b' = 0$ be the equation of B , with $A' = a'_1 + \dots + a'_n \neq 0$ the sum of its partial derivatives. We obtain the following equalities:

$$\begin{aligned} d_{\nu,B} &= -\frac{a'_1 \cdot \nu(x_1) + \dots + a'_n \cdot \nu(x_n) + b'}{A'} \\ \llbracket F(c) \rrbracket(\nu + d_{\nu,B}) &= \frac{a_1 \cdot \nu(x_1) + \dots + a_n \cdot \nu(x_n) + A d_{\nu,B} + b}{a_y} \\ \llbracket F(c) \rrbracket(\nu + d_{\nu,B}) \cdot A' a_y &= (A' a_1 - A a'_1) \cdot \nu(x_1) + \dots + (A' a_n - A a'_n) \cdot \nu(x_n) + (A' b - A b') \\ \text{Pre}_{e,c,B}(F)(\nu) \cdot A' a_y &= \sum_{i=1}^n (A' a_i - A a'_i - a_y a'_i \cdot \text{wt}(\ell)) \cdot \nu(x_i) \\ &\quad + (A' b - A b' - a_y b' \cdot \text{wt}(\ell) + A' a_y \cdot \text{wt}(e)) \end{aligned}$$

Thus $\text{Pre}_{e,c,B}(F)(\nu)$ is described by an equation $a_y^f y = a_1^f x_1 + \dots + a_n^f x_n + b^f$. The announced bound on constants holds for $n \geq 1$. \square

Let us now bring everything together for the value iteration operator \mathcal{F} .

Proposition 7.9. *Let $\mathbf{V} = (P_\ell, F_\ell)_{\ell \in L}$ be a piecewise affine value function, where every (P_ℓ, F_ℓ) has complexity at most $\langle m, \beta \rangle$, Let $q = |E|$ be the number of edges in \mathcal{G} . Then there exists a piecewise affine value function $\mathbf{V}' = (P'_\ell, F'_\ell)_{\ell \in L}$ such that $\mathbf{V}' = \mathcal{F}(\mathbf{V})$. In addition, every (P'_ℓ, F'_ℓ) has complexity at most $\langle m', \beta' \rangle$ where $m' = 36q^2(4m + 8n + 6)^{3n(n+2)}$ and β' is polynomial in n , $\max(w_{\max}^L, w_{\max}^E)$ and β .*

Proof. Fix a location $\ell \in L_{\text{Min}}$. The case L_{Max} is symmetrical and will not be detailed. If $\ell \in L_t$, let $(P'_\ell, F'_\ell) = (P_\ell, F_\ell)$ of complexity at most $\langle m, \beta \rangle$. Otherwise,

$$\mathbf{V}'_\ell = \min_{e=(\ell,g,\mathcal{Y},\ell')} \text{Pre}_e(\text{Guard}_g(\text{Unreset}_\mathcal{Y}(\mathbf{V}'_{\ell'}))).$$

For every edge $e = (\ell, g, \mathcal{Y}, \ell')$, we construct a piecewise affine value function (P_e, F_e) encoding $\text{Pre}_e(\text{Guard}_g(\text{Unreset}_\mathcal{Y}(\mathbf{V}'_{\ell'})))$.

We construct from (P_ℓ, F_ℓ) , by Lemmas 7.5, 7.6 and 7.7, an atomic tube partition (P_1, F_1) encoding $\text{Guard}_g(\text{Unreset}_\mathcal{Y}(\mathbf{V}'_{\ell'}))$. Following complexity results given in these lemmas, we end up with a complexity at most $\langle m_1, \beta_1 \rangle$, such that:

$$\begin{cases} m_1 &= (m + 2n)(1 + (m + 2n)\text{Splits}(m + 2n, n)) \\ \beta_1 &= 2n \max(\beta, M)^2 \end{cases}$$

As explained before (see equation 7.3), $\text{Pre}_e(\llbracket F_1 \rrbracket)$ is obtained using computations on base cells of P_1 followed by a minimum with $\llbracket F_1 \rrbracket$. Given a base cell c_b of P_1 , a cell c in \mathcal{C}_{c_b} , and a border B in $\mathcal{B}_{c_b}(c)$, we can apply Lemma 7.8 to deduce the existence of an affine value function $f_{e,c,B}$ such that $f_{e,c,B} = \text{Pre}_{e,c,B}(F_1)$ on c_b . We then have, for ν in base cell c_b :

$$\begin{aligned} \text{Pre}_e(\llbracket F_1 \rrbracket)(\nu) &= \min(\llbracket F_1 \rrbracket(\nu), \min_{c \in \mathcal{C}_{c_b}} \min_{B \in \mathcal{B}_{c_b}(c)} \text{Pre}_{e,c,B}(F_1)(\nu)) \\ &= \min(\llbracket F_1 \rrbracket(\nu), \min_{c \in \mathcal{C}_{c_b}} \min_{B \in \mathcal{B}_{c_b}(c)} f_{e,c,B}(\nu)) \end{aligned}$$

More precisely, for each base cell c_b of P_1 , we can compute the affine value function $f_{e,c,B}$ over c_b , with Lemma 7.8, with coefficients bounded by $\beta_2 = 4n\beta_1^2 \max(w_{\max}^L, w_{\max}^E)$. There are at most $\text{Splits}(m, n)$ cells in \mathcal{C}_{c_b} and at most 2 borders in $\mathcal{B}_{c_b}(e)$, thus the above formula for $\text{Pre}_e(\llbracket F_1 \rrbracket)$ involves a maximum amongst $q_1 = 2\text{Splits}(m_1, n) + 1$ elements. Applying Lemma 7.4, we can thus compute a piecewise affine value function (P_{c_b}, F_{c_b}) representing $\text{Pre}_e(\llbracket F_1 \rrbracket)$ over c_b . Using complexity results given in Lemma 7.4, we can deduce that constant complexity is bounded by $\beta_3 = 2\beta_2^2$, and that size complexity is bounded by $m_3 = q_1 m_1 + q_1^2$.

Then, we can use all (P_{c_b}, F_{c_b}) to define a unique (P_e, F_e) , by merging all P_{c_b} using the operator \oplus and by setting $\llbracket F_e \rrbracket(\nu) = \llbracket F_{c_b} \rrbracket(\nu)$ with $c_b = [\nu]_{P_1}$. The resulting size complexity m_e thus satisfies $m_e \leq m_3 \text{Splits}(m_1, n)$, while the constant complexity is left unchanged with $\beta_e = \beta_3$.

Finally, we can apply Lemma 7.4 to compute $\min_{e=(\ell, g, \mathcal{Y}, \ell')} \llbracket F_e \rrbracket$. The result is a partition function (P'_ℓ, F'_ℓ) encoding \mathbf{V}'_ℓ . Regarding complexity, Lemma 7.4 ensure that the value function \mathbf{V}'_ℓ has complexity at most $\langle m', \beta' \rangle$, with:

$$m' = qm_e + q^2 \quad \text{and} \quad \beta' = 2\beta_e^2$$

Putting everything together, we obtain :

$$\begin{aligned} \text{Splits}(m, n) &\leq (2m + 2)^n && \text{(Lemma 2.6)} \\ m_1 &\leq (2m + 4n + 2)^{n+2} && (1 \leq \text{Splits}(m + 2n, n)) \\ \text{Splits}(m_1, n) &\leq (4m + 8n + 6)^{n(n+2)} && ((2x^{n+2} + 2)^n \leq (2x + 2)^{n(n+2)}) \\ q_1 &\leq 3(4m + 8n + 6)^{n(n+2)} \\ m_3 &\leq 18(4m + 8n + 6)^{2n(n+2)} && (q_1 m_1 + q_1^2 \leq 2q_1^2) \\ m_e &\leq 18(4m + 8n + 6)^{3n(n+2)} \\ m' &\leq 36q^2(4m + 8n + 6)^{3n(n+2)} && (qm_e + q^2 \leq 2q^2 m_e) \end{aligned}$$

Regarding constant complexity, we have:

$$\begin{aligned} \beta' &= 2(\beta_e)^2 = 2(\beta_3)^2 = 2^3 \beta_2^4 \\ &= 2^3 (4n\beta_1^2 \max(w_{\max}^L, w_{\max}^E))^4 && (\beta_2 = 4n\beta_1^2 \max(w_{\max}^L, w_{\max}^E)) \\ &= 2^{11} \max(w_{\max}^L, w_{\max}^E)^4 n^4 (2n \max(\beta, M)^2)^8 && (\beta_1 = 2n \max(\beta, M)^2) \\ &= 2^{19} \max(w_{\max}^L, w_{\max}^E)^4 n^{12} \max(\beta, M)^{16} \quad \square \end{aligned}$$

Now that we know how to perform one step of computation of \mathcal{F} , we can estimate the complexity of iterated computations of this operator.

Corollary 7.10. *Let $\mathbf{V} = (P_\ell, F_\ell)_{\ell \in L}$ be a piecewise affine value function, where every (P_ℓ, F_ℓ) has complexity at most $\langle m, \beta \rangle$. Let $q = |E|$ be the number of edges in \mathcal{G} . For every $i \geq 1$, we can compute a piecewise affine value function $\mathbf{V}^{(i)} = (P_\ell^{(i)}, F_\ell^{(i)})_{\ell \in L}$ such that $\mathbf{V}^{(i)} = \mathcal{F}^i(\mathbf{V})$. In addition, every $(P_\ell^{(i)}, F_\ell^{(i)})$ has complexity at most $\langle m^{(i)}, \beta^{(i)} \rangle$ where $m^{(i)} = \max(m, 8n + 6, 10q)^{(6n(n+2))^i}$ and $\beta^{(i)}$ is polynomial in n , $\max(w_{\max}^L, w_{\max}^E)$ and β .*

Proof. The result can be obtained by induction on i . The initialisation is $m \leq m^{(0)}$. For the induction step, by Proposition 7.9 we must show $36q^2(4m^{(i)} + 8n + 6)^{3n(n+2)} \leq m^{(i+1)}$. Observe that $36q^2 \leq (2q)^{3n(n+2)}$, and $8n + 6, 10q \leq m^{(i)}$, so that $36q^2(4m^{(i)} + 8n + 6)^{3n(n+2)} \leq (10qm^{(i)})^{3n(n+2)} \leq (m^{(i)})^{6n(n+2)} \leq m^{(i+1)}$. This bound is crude but sufficient for our results.

We note that even finer bounds on the iteration of $m \mapsto 36q^2(4m + 8n + 6)^{3n(n+2)}$ lead to doubly-exponential complexities. \square

Thanks to Lemma 7.3, it will follow that \mathbf{V}^i , obtained by applying i times the operator \mathcal{F} over \mathbf{V}^0 of bounded complexity, can be stored using a space depending polynomially on m_0 and β_0 , exponentially on n , and doubly-exponentially on i . Moreover, the operations that we have described can be performed with a time complexity at most linear in the size of the output. This provides a doubly-exponential upper bound on the complexity (with respect to i) of computing \mathbf{V}^i . This concludes the proof of Theorem 7.1.

7.4. Exponential vs doubly-exponential. Let us now discuss the result of [ABM04], where an exponential upper bound on the bounded value problem is obtained with a non-symbolic algorithm on a slightly different setting (with non-negative weights only, without final weights, in a concurrent setting).

Whereas the concurrent setting they use generalizes ours, the sign of weights has seemingly no impact in the proofs, and the symbolic version requires minor changes related to the continuity of value functions and to the way guards are handled. These changes should not affect the complexity significantly. The main difference with their work is that their partition object is nested, it forms a tree structure where cells are partitionned into sub-cells as needed. Our techniques can be extended to nested partitions in the same way, and this has been detailed in [BG19, Chapter 10]. However, even with nested partitions the complexity of our techniques is still doubly-exponential, and the reason for this exponential gap is not apparent.

As a tentative answer, we make the following observation. If the game has no resets, *i.e.* $\mathcal{Y} = \emptyset$ on all edges, the complexity of our approach becomes exponential.¹⁸ In [ABM04], the way one should deal with resets is not detailed, it is therefore left open whether we could obtain an exponential bound or whether their solution is in fact doubly-exponential.

As an additional point of interest, it is shown in [BG19, Chapter 10.1.4] that there exists a tube partition P , of complexity m , such that if P' of complexity m' is obtained after applying \mathcal{F} , then $m' = \Theta(m^{n-1})$. This is the root of the issue, as we would need $m' = \mathcal{O}(m)$ in order to obtain an exponential bound when nesting \mathcal{F} .

7.5. Bounding partial derivatives. In the previous analysis, we explained that constants (partial derivatives and additive constants) grow polynomially at each elementary step, which is enough for an exponential upper bound (double-exponential growth of their value, stored in binary). This rough analysis will not be fine enough for some of our results. In particular, the approximation results of Section 8 will be sensitive to the partial derivatives in a linear (and not logarithmic) way. In this section, we study the growth of these partial derivatives more closely. This time, our focus will not be on the space required to store affine equations, but rather on mathematical properties of the value functions, namely their Lipschitz-continuity, closely related to bounds on partial derivatives. As a result, we revert to denoting affine equations as terms $y = a_1x_1 + \dots + a_nx_n + b$ with rational constants instead of using integers with a separately stored denominator a_y .

¹⁸This requires nested partitions and a more involved complexity analysis detailed in [BG19, Chapter 10]. In this case, the `Unresety` steps can be skipped, and in Proposition 7.9 we end up with a *linear* growth of the complexity with respect to m , instead of *polynomial*.

Definition 7.11. A value function $\mathbf{V} : L \times \mathbb{R}_{\geq 0}^{\mathcal{X}} \rightarrow \mathbb{R}_{\infty}$ is said to be Λ -Lipschitz-continuous on regions, with $\Lambda \in \mathbb{R}_{\geq 0}$ when, for all regions r and all $\ell \in L$, $|\mathbf{V}(\ell, \nu) - \mathbf{V}(\ell, \nu')| \leq \Lambda \|\nu - \nu'\|_{\infty}$ for all valuations $\nu, \nu' \in r$, where $\|\nu\|_{\infty} = \max_{x \in \mathcal{X}} |\nu(x)|$ is the ∞ -norm of vector $\nu \in \mathbb{R}^{\mathcal{X}}$. The function \mathbf{V} is said to be Lipschitz-continuous on regions if it is Λ -Lipschitz-continuous on regions, for some $\Lambda \in \mathbb{R}_{\geq 0}$.

Since final weight functions fwt are piecewise affine and continuous on regions, they are Lipschitz-continuous on regions, so is \mathbf{V}^0 . We will maintain as an invariant that \mathbf{V}^i is Lipschitz-continuous on regions:

Proposition 7.12. *If every final weight in a WTG \mathcal{G} is Λ -Lipschitz-continuous on regions (and piecewise affine), and $i > 0$, then $\text{Val}_{\mathcal{G}}^i$ is Λ' -Lipschitz-continuous on regions, with Λ' polynomial in Λ , w_{\max}^L and n , and exponential in i . More precisely, Λ' is bounded by*

$$\frac{\max(n-1, 2)^i - 1}{\max(n-2, 1)} w_{\max}^L + (n-1)^i \max(\Lambda, w_{\max}^L).$$

Note that for a piecewise affine function with finitely many pieces, being Λ -Lipschitz-continuous on regions is equivalent to being continuous on regions and having all partial derivatives bounded by Λ in absolute value. The rest of this section is dedicated to proving Proposition 7.12, as a corollary of the following result.

Lemma 7.13. *If for all $\ell \in L$, \mathbf{V}_{ℓ} is piecewise affine with finitely many pieces that have all their partial derivatives bounded by Λ in absolute value, then for all $\ell \in L$, $\mathcal{F}(\mathbf{V})_{\ell}$ is continuous on regions and piecewise affine with partial derivatives bounded by $\max(\Lambda, |\text{wt}(\ell)| + (n-1)\Lambda)$ in absolute value.*

Proof. We will show that for every region r , $\mathcal{F}(\mathbf{V})$ restricted to r has those properties. Note that they are transmitted over finite min and max operations. The continuity on regions is easy to prove because it is stable by infimum and supremum. There exists a partition function (P_{ℓ}, F_{ℓ}) for each $\ell \in L$ that represents \mathbf{V} . As explained before, a crucial property is that, for a given valuation ν , the delays d that need to be considered in the infimum and supremum of $\mathcal{F}(\mathbf{V})_{\ell}(\nu)$ correspond to the intersection points of the diagonal half line containing the time successors of ν and borders of cells (if ν^b is such a valuation, $d = \|\nu^b - \nu\|_{\infty}$ is the associated delay). In particular, there is a finite number of such borders, and the final $\mathcal{F}(\mathbf{V})_{\ell}$ function can be written as a finite nesting of finite min and max operations over affine terms, each corresponding to a choice of delay and an edge to take. Formally, there are several cases to consider to define those terms, depending on delay and edge choices. For each available edge e , those terms can either be:

- If a delay 0 is taken and all clocks in $\mathcal{Y} \subseteq \mathcal{X}$ are reset by e , then

$$\text{wt}_{\Sigma}((\ell, \nu) \xrightarrow{0} (\ell, \nu) \xrightarrow{e} (\ell', \nu[\mathcal{Y} := 0])) = \text{wt}_{\Sigma}(e) + \mathbf{V}_{\ell'}(\nu[\mathcal{Y} := 0])$$

- If a delay $d > 0$ (leading to valuation ν^b on border B) is taken and all clocks in $\mathcal{Y} \subseteq \mathcal{X}$ are reset by e , then

$$\text{wt}_{\Sigma}((\ell, \nu) \xrightarrow{d} (\ell, \nu^b) \xrightarrow{e} (\ell', \nu^b[\mathcal{Y} := 0])) = \text{wt}_{\Sigma}(\ell) \cdot d + \text{wt}_{\Sigma}(e) + \mathbf{V}_{\ell'}(\nu^b[\mathcal{Y} := 0])$$

In the first case, the resulting partial derivatives are 0 for clocks in \mathcal{Y} , and the same as the partial derivatives in $\mathbf{V}_{\ell'}$ for all other clocks, which allows us to conclude that they are bounded by Λ . We now consider the second case. We argue that it can be decomposed as a

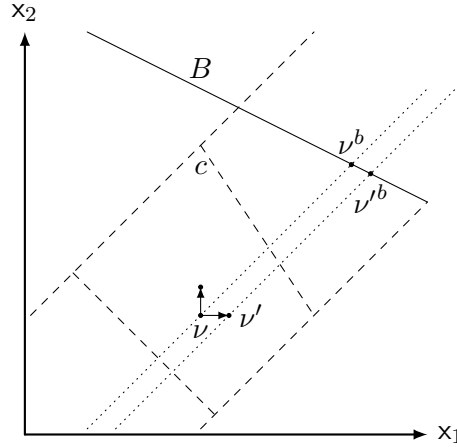


Figure 17: A cell c as described in the proof of Lemma 7.13. Dashed lines are borders of c , dotted lines are proof constructions.

delay followed by an edge of the first case, meaning that we can assume $\mathcal{Y} = \emptyset$ without loss of generality.

There are again two cases: the border B being inside a region or on the frontier of a region.

If the border is not the frontier of a region, it is the intersection points of two affine pieces of \mathbf{V}_{ℓ} whose equations (in the space \mathbb{R}^{n+1} whose n first coordinates are the clocks (x_1, \dots, x_n) and the last coordinate y corresponds to the value $\mathbf{V}_{\ell}(x_1, \dots, x_n)$) can be written $y = a_1x_1 + \dots + a_nx_n + b$ (before the border) and $y = a'_1x_1 + \dots + a'_nx_n + b'$ (after the border). Therefore, valuations at the borders all fulfill the equation

$$(a'_1 - a_1)x_1 + \dots + (a'_n - a_n)x_n + b' - b = 0 \quad (7.4)$$

We let $A = (a'_1 - a_1) + \dots + (a'_n - a_n)$. Consider that ℓ is a location of **Min** (the very same reasoning applies to the case of a location of **Max**). Since \mathcal{F} computes an infimum, we know that the function mapping the delay d to the weight obtained from reaching $\nu + d$ is decreasing before the border and increasing after. These functions are locally affine which implies that their slopes satisfy:

$$\text{wt}(\ell) + a_1 + \dots + a_n \leq 0 \quad \text{and} \quad \text{wt}(\ell) + a'_1 + \dots + a'_n \geq 0 \quad (7.5)$$

We deduce from these two inequalities that $A \geq 0$. The case where $A = 0$ would correspond to the case where the border contains a diagonal line, which is forbidden, and thus $A > 0$. Consider now a valuation ν and another valuation ν' of coordinates $(\nu(x_1), \dots, \nu(x_{k-1}), \nu(x_k) + \lambda, \nu(x_{k+1}), \dots, \nu(x_n))$ with λ small enough. The delays d and d' needed to arrive to the border starting from these two valuations are such that $\nu + d$ and $\nu' + d'$ both satisfy (7.4). We can then deduce

$$d' - d = \lambda \frac{a_k - a'_k}{A}$$

It is now possible to compute the partial derivative of $\mathcal{F}(\mathbf{V})_{\ell}$ in the k -th coordinate using the limit when λ tends to 0 of the quotient

$$\frac{\mathcal{F}(\mathbf{V})_{\ell, \nu'} - \mathcal{F}(\mathbf{V})_{\ell, \nu}}{\lambda} = \frac{\text{wt}(\ell)(d' - d) + \mathbf{V}_{\ell, \nu' + d'} - \mathbf{V}_{\ell, \nu + d}}{\lambda} \quad (7.6)$$

We may compute it by using the equations of the affine pieces before or after the border. We thus obtain

$$\begin{aligned}\frac{\mathcal{F}(\mathbf{V})_{\ell,\nu'} - \mathcal{F}(\mathbf{V})_{\ell,\nu}}{\lambda} &= \frac{a_k - a'_k}{A} (\text{wt}(\ell) + a_1 + \cdots + a_n) + a_k \\ \frac{\mathcal{F}(\mathbf{V})_{\ell,\nu'} - \mathcal{F}(\mathbf{V})_{\ell,\nu}}{\lambda} &= \frac{a_k - a'_k}{A} (\text{wt}(\ell) + a'_1 + \cdots + a'_n) + a'_k\end{aligned}$$

In the case where $a_k \geq a'_k$, the first equation, with (7.5), allows us to obtain that the quotient (and thus the partial derivative) is at most a_k . The second equation, with (7.5), allows us to obtain that the partial derivative is at least a'_k . In case $a'_k \geq a_k$, we obtain similarly that the partial derivative is at most a'_k and at least a_k . Since a_k and a'_k are bounded in absolute value by Λ , so is the partial derivative.

We now come back to the case where the border is on the frontier of a region. Then, it is a segment of a line of equation $x_k = c$ for some k and c . Moreover, $\mathbf{V}_{\ell'}$ contains at most three values for points of B : the limit coming from before the border, the value at the border, and the limit coming from after the border. The computation of $\mathcal{F}(\mathbf{V})$ considers values obtained from all three and takes the minimum (or maximum).

Now, let $y = a_1x_1 + \cdots + a_nx_n + b$ be the equation defining the affine piece of $\mathbf{V}_{\ell'}$ before the border (respectively at the border, after the border). Consider a valuation ν and another valuation ν' of coordinates $(\nu(x_1), \dots, \nu(x_{j-1}), \nu(x_j) + \lambda, \nu(x_{j+1}), \dots, \nu(x_n))$ with λ small enough. The delays d and d' needed to arrive to the border starting from these two valuations are such that $\nu + d$ and $\nu' + d'$ both satisfy $x_k = c$. We can then deduce that $d' - d = 0$ if $j \neq k$ and $d' - d = -\lambda$ if $j = k$. It is now possible to compute the partial derivative of $\mathcal{F}(\mathbf{V})_{\ell}$ in the j -th coordinate using (7.6) again. We may compute it by using the equations of the affine piece before the border (respectively at the border, after the border). Then,

$$\begin{aligned}\mathbf{V}_{\ell',\nu+d} &= a_1(x_1 + d) + \cdots + a_n(x_n + d) + b = \sum_{i=1, i \neq k}^n a_i(x_i + d) + a_k c + b \\ \mathbf{V}_{\ell',\nu'+d'} &= \sum_{i=1, i \neq k}^n a_i(x_i + d') + a_k c + b\end{aligned}$$

We thus obtain

$$\frac{\mathcal{F}(\mathbf{V})_{\ell,\nu'} - \mathcal{F}(\mathbf{V})_{\ell,\nu}}{\lambda} = \begin{cases} a_j & \text{if } j \neq k \\ -\text{wt}(\ell) - \sum_{i=1, i \neq k}^n a_i & \text{otherwise} \end{cases}$$

Then, the partial derivatives are bounded, in absolute value, by $|\text{wt}(\ell)| + (n - 1)\Lambda$. \square

As a corollary, we obtain Proposition 7.12 by a simple induction on i .

8. COMPUTING VALUES

In this section we conclude the proofs of Theorems 3.4 and 3.9, with a triply-exponential upper bound on computing (respectively approximating) values in divergent (respectively almost-divergent) weighted timed games.

These upper bounds are obtained by computations performed on the semi-unfolding $\mathcal{T}(\mathcal{G})$ described in Section 6. Whereas the computation is direct for divergent WTGs, we then extend it to almost-divergent WTGs by first explaining how to compute value approximations in kernels.

8.1. Divergent WTGs. We first explain how to compute the values of a divergent WTG \mathcal{G} , thus proving Theorem 3.4. By definition, in such a divergent WTG, there are no 0-cycles, and thus the kernel \mathbf{K} is empty. In this case, the semi-unfolding $\mathcal{T}(\mathcal{G})$ is a tree of depth i exponential in \mathcal{G} , and thus of doubly-exponential size. Proposition 6.1 ensures that the values of $\mathcal{T}(\mathcal{G})$ coincide with the values of \mathcal{G} . By Theorem 7.1, we can compute the (exact) values in $\mathcal{T}(\mathcal{G})$ in time doubly-exponential in i and exponential in the size of $\mathcal{T}(\mathcal{G})$. We thus obtain a triply-exponential algorithm computing the value of a divergent WTG. Equivalently, this shows that the value problem is in 3-EXPTIME for divergent WTGs.

8.2. Approximation of kernels. In order to generalise our computations from divergent to almost-divergent WTGs, we start by approximating a kernel \mathbf{K} of a WTG \mathcal{G} by extending the region-based approximation schema of [BJM15]. More precisely, we thus consider here a WTG \mathcal{G} composed of a kernel (a subgraph of a region game containing only 0-cycles, as defined in Section 5.2) and some target locations equipped with final weights. In the setting of [BJM15], not only cycles but all finite plays in kernels have weight 0, allowing for a reduction to a finite game. In our setting, we have to approximate the timed dynamics of runs, and therefore resort to the corner-point abstraction (as shown in the right part of Figure 11).

Our goal is to compute an ε -approximation of the value of the kernel (in a given initial configuration). Since final weight functions are piecewise affine with a finite number of pieces and continuous on regions, they are Λ -Lipschitz-continuous on regions, for a given constant $\Lambda \geq 0$. The technique to obtain the approximation is to consider regions of a refined enough granularity: we thus pick

$$N(\varepsilon, \Lambda) = \left\lceil \frac{w_{\max}^L |\mathcal{R}(\mathcal{G})| + \Lambda}{\varepsilon} \right\rceil \quad (8.1)$$

later denoted N when the parameters ε and Λ are clear from context.

Consider then the corner-point abstraction $\Gamma_N(\mathcal{G})$ described in Section 2.4, with locations of the form (ℓ, r, \mathbf{v}) such that \mathbf{v} is a corner of the $1/N$ -region r . Two plays ρ of \mathcal{G} and ρ' of $\Gamma_N(\mathcal{G})$ are said to be $1/N$ -close if they follow the same path \mathbf{p} in $\mathcal{R}_N(\mathcal{G})$. In particular, at each step the configurations (ℓ, ν) in ρ and (ℓ', r', ν') in ρ' (with ν' a corner of the $1/N$ -region r') satisfy $\ell = \ell'$ and $\nu \in r'$, and the edges taken in both plays have the same discrete weights. Close plays have *close* weights, in the following sense:

Lemma 8.1. *For all $1/N$ -close plays ρ of \mathcal{G} and ρ' of $\Gamma_N(\mathcal{G})$,*

$$|\text{wt}_{\mathcal{G}}(\rho) - \text{wt}_{\Gamma_N(\mathcal{G})}(\rho')| \leq \varepsilon$$

Proof. Since ρ and ρ' encounter the same locations of \mathcal{G} , one reaches a target location if and only if the other does. In the case where they do not reach a target location, both weights are infinite, and thus equal. We now look at the case where both plays reach a target location, moreover in the same step.

Consider the region path \mathbf{p} of the run ρ : it can be decomposed into a simple path with maximal cycles in it. The number of such maximal cycles is bounded by $|\mathcal{R}(\mathcal{G})|$ and the remaining simple path has length at most $|\mathcal{R}(\mathcal{G})|$. Since all cycles of a kernel are 0-cycles, the parts of ρ that follow the maximal cycles have weight exactly 0.

Consider the same decomposition for the play ρ' . Cycles of \mathbf{p} do not necessarily map to cycles over locations of $\Gamma_N(\mathcal{G})$, since the $1/N$ -corners could be distinct. However, Lemma 2.16 shows that, for all those cycles of \mathbf{p} , there exists a sequence of finite plays of \mathcal{G} whose weight

tends to the weight of ρ' . Since all those finite plays follow a cycle of the region game $\mathcal{R}(\mathcal{G})$ (with \mathcal{G} being a kernel), they all have weight 0. Hence, the parts of ρ' that follow the maximal cycles of \mathbf{p} have also weight exactly 0.

Therefore, the difference $|\mathbf{wt}_{\mathcal{G}}(\rho) - \mathbf{wt}_{\Gamma_N(\mathcal{G})}(\rho')|$ is concentrated on the remaining simple path of \mathbf{p} : on each edge of this path, the maximal weight difference is $1/N \times w_{\max}^L$ since $1/N$ is the largest difference possible in time delays between plays that stay $1/N$ -close (since they stay in the same $1/N$ -regions). Moreover, the difference between the final weight functions is bounded by Λ/N , since the final weight function \mathbf{fwt} is Λ -Lipschitz-continuous and the final weight function of $\Gamma_N(\mathcal{G})$ is obtained as limit of \mathbf{fwt} . Summing the two contributions, we obtain as upper bound the constant $(w_{\max}^L |\mathcal{R}(\mathcal{G})| + \Lambda)/N \leq \varepsilon$. \square

In particular, if we start in some configurations (ℓ, ν) of \mathcal{G} , and $((\ell, r, \mathbf{v}), \mathbf{v})$ of $\Gamma_N(\mathcal{G})$, with $\nu \in r$, since both players have the ability to stay $1/N$ -close all along the plays, a bisimulation argument allows us to obtain that the values of the two games are also close in (ℓ, ν) and $((\ell, r, \mathbf{v}), \mathbf{v})$, respectively:

Lemma 8.2. *For all locations $\ell \in L$, $1/N$ -regions r , valuations $\nu \in r$ and corners \mathbf{v} of r ,*

$$|\mathbf{Val}_{\mathcal{G}}(\ell, \nu) - \mathbf{Val}_{\Gamma_N(\mathcal{G})}((\ell, r, \mathbf{v}), \mathbf{v})| \leq \varepsilon$$

Proof. To obtain the result, we prove that

$$\mathbf{Val}_{\mathcal{G}}(\ell, \nu) \leq \mathbf{Val}_{\Gamma_N(\mathcal{G})}((\ell, r, \mathbf{v}), \mathbf{v}) + \varepsilon \quad \text{and} \quad \mathbf{Val}_{\Gamma_N(\mathcal{G})}((\ell, r, \mathbf{v}), \mathbf{v}) \leq \mathbf{Val}_{\mathcal{G}}(\ell, \nu) + \varepsilon$$

By definition and determinacy of turn-based WTGs, this is equivalent to proving these two inequalities:

$$\begin{aligned} \inf_{\sigma_{\text{Min}}} \sup_{\sigma_{\text{Max}}} \mathbf{wt}_{\mathcal{G}}(\mathbf{play}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) &\leq \inf_{\sigma'_{\text{Min}}} \sup_{\sigma'_{\text{Max}}} \mathbf{wt}_{\Gamma_N(\mathcal{G})}(\mathbf{play}(((\ell, r, \mathbf{v}), \mathbf{v}), \sigma'_{\text{Min}}, \sigma'_{\text{Max}})) + \varepsilon \\ \sup_{\sigma'_{\text{Max}}} \inf_{\sigma'_{\text{Min}}} \mathbf{wt}_{\Gamma_N(\mathcal{G})}(\mathbf{play}(((\ell, r, \mathbf{v}), \mathbf{v}), \sigma'_{\text{Min}}, \sigma'_{\text{Max}})) &\leq \sup_{\sigma_{\text{Max}}} \inf_{\sigma_{\text{Min}}} \mathbf{wt}_{\mathcal{G}}(\mathbf{play}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) + \varepsilon \end{aligned}$$

To show the first inequality, it suffices to show that for all σ'_{Min} , there exists σ_{Min} such that for all σ_{Max} , there is σ'_{Max} verifying

$$|\mathbf{wt}_{\mathcal{G}}(\mathbf{play}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) - \mathbf{wt}_{\Gamma_N(\mathcal{G})}(\mathbf{play}(((\ell, r, \mathbf{v}), \mathbf{v}), \sigma'_{\text{Min}}, \sigma'_{\text{Max}}))| \leq \varepsilon \quad (8.2)$$

For the second, it suffices to show that for all σ'_{Max} , there exists σ_{Max} such that for all σ_{Min} , there is σ'_{Min} verifying the same equation (8.2). We will detail the proof for the first, the second being syntactically the same, with both players swapped.

Equation (8.2) can be obtained from Lemma 8.1, under the condition that the plays $\mathbf{play}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})$ and $\mathbf{play}(((\ell, r, \mathbf{v}), \mathbf{v}), \sigma'_{\text{Min}}, \sigma'_{\text{Max}})$ are $1/N$ -close. Therefore, we fix a strategy σ'_{Min} of Min in the game $\Gamma_N(\mathcal{G})$, and we construct a strategy σ_{Min} of Min in \mathcal{G} , as well as two mappings $f: \mathbf{FPlays}_{\mathcal{G}}^{\text{Min}} \rightarrow \mathbf{FPlays}_{\Gamma_N(\mathcal{G})}^{\text{Min}}$ and $g: \mathbf{FPlays}_{\Gamma_N(\mathcal{G})}^{\text{Max}} \rightarrow \mathbf{FPlays}_{\mathcal{G}}^{\text{Max}}$ such that:

- for all $\rho \in \mathbf{FPlays}_{\mathcal{G}}^{\text{Min}}$, ρ and $f(\rho)$ are $1/N$ -close, and if ρ is consistent with σ_{Min} and starts in (ℓ, ν) , then $f(\rho)$ is consistent with σ'_{Min} and starts in $((\ell, r, \mathbf{v}), \mathbf{v})$;
- for all $\rho' \in \mathbf{FPlays}_{\Gamma_N(\mathcal{G})}^{\text{Max}}$, $g(\rho')$ and ρ' are $1/N$ -close, and if ρ' is consistent with σ'_{Min} and starts in $((\ell, r, \mathbf{v}), \mathbf{v})$, then $g(\rho')$ is consistent with σ_{Min} and starts in (ℓ, ν) .

We build σ_{Min} , f , and g by induction on the length k of plays, over prefixes of plays of length $k-1$, k and k , respectively. For $k=0$ (plays of length 0 are those restricted to a single configuration), we let $f(\ell, \nu) = ((\ell, r, \mathbf{v}), \mathbf{v})$ and $g((\ell, r, \mathbf{v}), \mathbf{v}) = (\ell, \nu)$, leaving the other values arbitrary (since we will not use them).

Then, we suppose σ_{Min} , f , and g built until length $k - 1$, k and k , respectively (if $k = 0$, σ_{Min} has not been build yet), and we define them on plays of length k , $k + 1$ and $k + 1$, respectively. For every $\rho \in \text{FPlays}_{\mathcal{G}}^{\text{Min}}$ of length k , we note $\rho' = f(\rho)$. Consider the decision $(d', e') = \sigma'_{\text{Min}}(\rho')$ and ρ'_+ the prefix ρ' extended with the decision (d', e') . By timed bisimulation, there exists (d, e) such that the prefix ρ_+ composed of ρ extended with the decision (d, e) builds $1/N$ -close plays ρ_+ and ρ'_+ . We let $\sigma_{\text{Min}}(\rho) = (d, e)$. If $\rho_+ \in \text{FPlays}_{\mathcal{G}}^{\text{Min}}$, we also let $f(\rho_+) = \rho'_+$, and otherwise we let $g(\rho'_+) = \rho_+$. Symmetrically, consider $\rho' \in \text{FPlays}_{\Gamma_N(\mathcal{G})}^{\text{Max}}$ of length k , and $\rho = g(\rho')$. For all possible decisions (d', e') , by timed bisimulation, there exists a decision (d, e) in the prefix ρ such that the respective extended plays ρ'_+ and ρ_+ are $1/N$ -close. We then let $g(\rho'_+) = \rho_+$ if $\rho_+ \in \text{FPlays}_{\mathcal{G}}^{\text{Max}}$ and $f(\rho_+) = \rho'_+$ otherwise. We extend the definition of f and g arbitrarily for other prefixes of plays. The properties above are then trivially verified.

We then fix a strategy σ_{Max} of Max in the game \mathcal{G} , which determines a unique play $\text{play}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})$. We construct a strategy σ'_{Max} of Max in the game $\Gamma_N(\mathcal{G})$ by building the unique play $\text{play}(((\ell, r, \nu), \nu), \sigma'_{\text{Min}}, \sigma'_{\text{Max}})$ we will be interested in, such that each of its prefixes is in relation, via f or g , to the associated prefix of $\text{play}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})$. Thus, we only need to consider a prefix of play $\rho' \in \text{FPlays}_{\Gamma_N(\mathcal{G})}^{\text{Max}}$ that starts in $((\ell, r, \nu), \nu)$ and is consistent with σ'_{Min} , and σ'_{Max} built so far. Consider the play $\rho = g(\rho')$, starting in (ℓ, ν) and consistent with σ_{Min} , and σ_{Max} (by assumption). For the decision $(d, e) = \sigma_{\text{Max}}(\rho)$ (letting ρ_+ be the extended prefix), the definition of f and g ensures that there exists a decision (d', e') after ρ' that results in an extended play ρ'_+ that is $1/N$ -close, via f or g , with ρ_+ . We thus can choose $\sigma'_{\text{Max}}(\rho') = (d', e')$.

We finally have built two plays $\text{play}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})$ and $\text{play}((\ell', \nu'), \sigma'_{\text{Min}}, \sigma'_{\text{Max}})$ that are $1/N$ -close, as needed, which concludes this proof. \square

We can thus obtain an ε -approximation of $\text{Val}_{\mathcal{G}}(\ell, \nu)$ by computing $\text{Val}_{\Gamma_N(\mathcal{G})}((\ell, r, \nu), \nu)$ for any corner ν of r . Recall that $\Gamma_N(\mathcal{G})$ can be considered as an untimed weighted game (with reachability objective). Thus we can apply the result of [BGHM16], where it is shown that the optimal values of such games can be computed in pseudo-polynomial time (*i.e.* polynomial with respect to the number of locations $|\Gamma_N(\mathcal{G})|$, at most $(n + 1)|\mathcal{R}_N(\mathcal{G})|$, and the weights of transitions w_{max} encoded in unary, instead of binary): more precisely, the value $\text{Val}_{\Gamma_N(\mathcal{G})}$ is obtained using the operator \mathcal{F} of (2.2) (page 20), as the i -th iteration with $i = ((2|\Gamma_N(\mathcal{G})| - 1)w_{\text{max}} + 1)|\Gamma_N(\mathcal{G})|$, polynomial in $|L|$, w_{max} , M and N , and exponential in n . We then define an ε -approximation of $\text{Val}_{\mathcal{G}}$, named Val'_N , on each $1/N$ -region by interpolating the values of its $1/N$ -corners in $\Gamma_N(\mathcal{G})$ with a piecewise affine function: If ν is a valuation that belongs to the $1/N$ -region r , then ν can be expressed as a (unique) convex combination of the $1/N$ -corners ν of r , so that $\nu = \sum_{\nu} \alpha_{\nu} \nu$ with $\alpha_{\nu} \in [0, 1]$ for all ν , and we let $\text{Val}'_N(\ell, \nu) = \sum_{\nu} \alpha_{\nu} \text{Val}_{\Gamma_N(\mathcal{G})}((\ell, r, \nu), \nu)$ for all locations ℓ of \mathcal{G} .

Moreover, we can control the growth of the Lipschitz constant of the approximated value for further use.

Lemma 8.3. *Val'_N is an ε -approximation of $\text{Val}_{\mathcal{G}}$, *i.e.* $\|\text{Val}'_N - \text{Val}_{\mathcal{G}}\|_{\infty} \leq \varepsilon$. Moreover, Val'_N is piecewise affine with a finite number of pieces and $2(w_{\text{max}}^L |\mathcal{R}(\mathcal{G})| + \Lambda)$ -Lipschitz-continuous over regions.*

Proof. By construction, the approximated value is piecewise affine with one piece per $1/N$ -region. To prove the Lipschitz constant, it is then sufficient to bound the difference between $\text{Val}_{\Gamma_N(\mathcal{G})}((\ell, r, \nu), \nu)$ and $\text{Val}_{\Gamma_N(\mathcal{G})}((\ell, r, \nu'), \nu')$, for ν and ν' two corners of a $1/N$ -region r . We

can pick any valuation ν in r and apply Lemma 8.2 twice, between ν and \mathbf{v} , and between ν and \mathbf{v}' . We obtain $|\text{Val}_{\Gamma_N(\mathcal{G})}((\ell, r, \mathbf{v}), \mathbf{v}) - \text{Val}_{\Gamma_N(\mathcal{G})}((\ell, r, \mathbf{v}'), \mathbf{v}')| \leq 2(w_{\max}^L |\mathcal{R}(\mathcal{G})| + \Lambda)/N = 2(w_{\max}^L |\mathcal{R}(\mathcal{G})| + \Lambda) \|\mathbf{v} - \mathbf{v}'\|_{\infty}$. \square

The time complexity needed to compute Val'_N is polynomial in $|L|$, w_{\max} , M and N , and exponential in n .

8.3. Approximation of almost-divergent WTGs. We now explain how to approximate the value of an almost-divergent WTG \mathcal{G} , thus proving Theorem 3.9. After having computed the semi-unfolding $\mathcal{T}(\mathcal{G})$ described in Section 6, we perform a bottom-up computation of the approximation. The addition of kernels (with respect to the case of divergent WTGs studied before) requires us to compute an approximation instead of the actual value. Notice that the techniques used in Theorem 7.1, applied for $i = 1$ step (and not for the whole tree as for divergent WTGs), allow us to compute the value of a non-kernel node of $\mathcal{T}(\mathcal{G})$, depending on the values of its children. There is no approximation needed here, so that if we have computed ε -approximations of all its children, we can compute an ε -approximation of the node. More formally, this is justified by the following lemma with $i = 1$:

Lemma 8.4. *Let \mathcal{G} and \mathcal{G}' be two WTGs that only differ on the final weight functions fwt and fwt' . Then, $\|\text{Val}_{\mathcal{G}} - \text{Val}_{\mathcal{G}'}\|_{\infty} \leq \|\text{fwt} - \text{fwt}'\|_{\infty}$. Moreover, for all $i \in \mathbb{N}$, $\|\text{Val}_{\mathcal{G}}^i - \text{Val}_{\mathcal{G}'}^i\|_{\infty} \leq \|\text{fwt} - \text{fwt}'\|_{\infty}$.*

Proof. Consider two strategies σ_{Min} and σ_{Max} for both players in \mathcal{G} and \mathcal{G}' (since both games are identical up-to the final weight functions, they share the same sets of strategies for both players). Then, from an initial configuration (ℓ, ν) , the plays $\text{play}_{\mathcal{G}}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})$ and $\text{play}_{\mathcal{G}'}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})$ are the same, and their weights only differ by the final weight functions taken at the same configurations. Thus,

$$|\text{wt}(\text{play}_{\mathcal{G}}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) - \text{wt}(\text{play}_{\mathcal{G}'}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}}))| \leq \|\text{fwt} - \text{fwt}'\|_{\infty}$$

In particular,

$$\text{wt}(\text{play}_{\mathcal{G}}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) \leq \text{wt}(\text{play}_{\mathcal{G}'}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) + \|\text{fwt} - \text{fwt}'\|_{\infty}$$

so that, taking infimum over all σ_{Min} and supremum over all σ_{Max} gives

$$\text{Val}_{\mathcal{G}}(\ell, \nu) \leq \text{Val}_{\mathcal{G}'}(\ell, \nu) + \|\text{fwt} - \text{fwt}'\|_{\infty}$$

A symmetrical argument allows us to conclude that $\|\text{Val}_{\mathcal{G}} - \text{Val}_{\mathcal{G}'}\|_{\infty} \leq \|\text{fwt} - \text{fwt}'\|_{\infty}$. Using the bounded horizon versions of Val and wt in the previous proof, we also obtain $\|\text{Val}_{\mathcal{G}}^i - \text{Val}_{\mathcal{G}'}^i\|_{\infty} \leq \|\text{fwt} - \text{fwt}'\|_{\infty}$ for any i . \square

We now explain in detail the full process of approximation of the value $\text{Val}_{\mathcal{G}}(\ell_0, \nu_0)$ of an almost-divergent WTG \mathcal{G} : it is a bottom-up computation on the tree T rooted in (ℓ_0, r_0) (with r_0 the region of ν_0) that we used to describe the semi-unfolding $\mathcal{T}(\mathcal{G})$. By Proposition 6.1, the value we want to approximate is equal to $\text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$. For a node s in T , let \mathbf{V}_s denote the exact value function of the corresponding node in $\mathcal{T}(\mathcal{G})$. In particular, the value function at the root of T is equal to $\nu \mapsto \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu)$. Our algorithm iteratively computes an approximated value function \mathbf{V}'_s for all nodes s of T .

To obtain an adequate ε -approximation of $\text{Val}_{\mathcal{G}}(\ell_0, \nu_0)$, we will thus need to guarantee a precision in kernels that depend on the number of kernels we visit in the semi-unfolding. Let α be the maximal number of kernels along a branch of the tree T . For a given node s in T ,

we also let $\alpha(s)$ be the maximal number of kernels along the branches of the subtree rooted in s . Finally, let us denote by $h(s)$ the maximal length of a branch in the subtree rooted in s .

We will maintain, along the bottom-up computation, that \mathbf{V}'_s is a Λ_s -Lipschitz-continuous mapping on regions such that

$$\Lambda_s \leq \max(2, n)^{h(s)} (2w_{\max}^L |\mathcal{R}(\mathcal{G})| + \Lambda) \quad \text{and} \quad \|\mathbf{V}_s - \mathbf{V}'_s\|_\infty \leq \alpha(s)\varepsilon/\alpha \quad (8.3)$$

where Λ is the maximal Lipschitz constant of the final weight functions of \mathcal{G} . In particular, at the root of T where $\alpha(s) = \alpha$, we indeed recover an ε -approximation of the value of the game in configuration (ℓ_0, ν_0) .

For each leaf node s of T , \mathbf{V}_s and \mathbf{V}'_s are equal to the final weight function fwt of $\mathcal{T}(\mathcal{G})$, and we thus get the invariant (8.3) for free (knowing that $\alpha(s) = h(s) = 0$).

For an internal node s of T (that either gives rise to a single state (ℓ, r) of $\mathcal{T}(\mathcal{G})$, or to a kernel $\mathbf{K}_{\ell, r}$), let us suppose that the value of each child s' of s in T have been computed and verify the invariant (8.3).

If s is a node of the form (ℓ, r) (that is not part of a kernel), we define two WTGs $\tilde{\mathcal{G}}$ and $\tilde{\mathcal{G}}'$ that contain the state (ℓ, r) as well as its children in T . The children s' are made target states with respective final weight functions given by $\mathbf{V}_{s'}$ and $\mathbf{V}'_{s'}$. Moreover, we know that

$$\alpha(s) = \max_{s'} \alpha(s') \quad \text{and} \quad h(s) = \max_{s'} h(s') + 1$$

By definition, \mathbf{V}_s is equal to $\nu \mapsto \text{Val}_{\tilde{\mathcal{G}}}(s, \nu)$ and thus to $\nu \mapsto \text{Val}_{\tilde{\mathcal{G}}}^1(s, \nu)$ since $\tilde{\mathcal{G}}$ is acyclic of depth 1. Our approximation algorithm consists in letting $\mathbf{V}'_s = \nu \mapsto \text{Val}_{\tilde{\mathcal{G}}'}^1(s, \nu)$. By Lemma 8.4 (with $i = 1$ and ε being $\max_{s'} \alpha(s')\varepsilon/\alpha$), $\|\mathbf{V}_s - \mathbf{V}'_s\|_\infty \leq \alpha(s)\varepsilon/\alpha$. By Lemma 7.13, \mathbf{V}'_s is Λ_s -Lipschitz-continuous on regions with $\Lambda_s \leq \max(\max_{s'} \Lambda_{s'}, |\text{wt}(\ell)| + (n-1) \max_{s'} \Lambda_{s'})$: with the help of the invariant (8.3) for all children s' , and since $|\text{wt}(\ell)| \leq w_{\max}^L$ and $h(s') \leq h(s) - 1$, we obtain $\Lambda_s \leq \max(2, n)^{h(s)} (2w_{\max}^L |\mathcal{R}(\mathcal{G})| + \Lambda)$.

If s is a node of the form $\mathbf{K}_{\ell, r}$, we define two WTGs $\tilde{\mathcal{G}}$ and $\tilde{\mathcal{G}}'$, that contain the locations of the kernel $\mathbf{K}_{\ell, r}$ of $\mathcal{T}(\mathcal{G})$, as well as the children of s in T (reached by output edges of $\mathbf{K}_{\ell, r}$). The children s' are made target states of respective final weight functions $\mathbf{V}_{s'}$ and $\mathbf{V}'_{s'}$. Moreover, we know that

$$\alpha(s) = \max_{s'} \alpha(s') + 1 \quad \text{and} \quad h(s) = \max_{s'} h(s') + 1$$

Thus, games $\tilde{\mathcal{G}}$ and $\tilde{\mathcal{G}}'$ are identical, except for their final weight functions that are ε -close. Thus, we know by Lemma 8.4 that $\|\text{Val}_{\tilde{\mathcal{G}}} - \text{Val}_{\tilde{\mathcal{G}}'}\|_\infty \leq \max_{s'} \alpha(s')\varepsilon/\alpha$. Moreover, by definition, \mathbf{V}_s is equal to $\nu \mapsto \text{Val}_{\tilde{\mathcal{G}}}(s, \nu)$. Our approximation algorithm consists in letting \mathbf{V}'_s be equal to an ε/α -approximation of $\text{Val}_{\tilde{\mathcal{G}}}$, obtained by Lemma 8.3 with a granularity $N(\varepsilon/\alpha, \max_{s'} \Lambda_{s'})$: we thus have $\|\text{Val}_{\tilde{\mathcal{G}}}(s, \cdot) - \mathbf{V}'_s\|_\infty \leq \varepsilon/\alpha$, and \mathbf{V}'_s is Λ_s -Lipschitz-continuous on regions with $\Lambda_s \leq 2(w_{\max}^L |\mathcal{R}(\mathcal{G})| + \max_{s'} \Lambda_{s'})$. By triangular inequality, we deduce that $\|\mathbf{V}_s - \mathbf{V}'_s\|_\infty \leq (\max_{s'} \alpha(s'))\varepsilon/\alpha + \varepsilon/\alpha = \alpha_s \varepsilon/\alpha$. Moreover, with the help of invariant (8.3) for $\Lambda_{s'}$, we once again obtain that $\Lambda_s \leq \max(2, n)^{h(s)} (2w_{\max}^L |\mathcal{R}(\mathcal{G})| + \Lambda)$.

We thus obtain an algorithm that faithfully computes an ε -approximation of the value of the game. Let us now discuss the complexity of the algorithm. Overall, the biggest Lipschitz constant for \mathbf{V}' is $\max(2, n)^h (2w_{\max}^L |\mathcal{R}(\mathcal{G})| + \Lambda)$ with h the height of the semi-unfolding that is $|\mathcal{R}(\mathcal{G})|(3|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 2) + 1$ as noticed in (6.1) (page 41). This Lipschitz constant is thus at most doubly-exponential with respect to the size of \mathcal{G} . Therefore, the biggest granularity N used in kernel approximations (described in (8.1)) can be globally

bounded as doubly-exponential in the size of \mathcal{G} and polynomial in $1/\varepsilon$. This entails that each kernel approximation can be performed in time doubly-exponential in the size of \mathcal{G} , and polynomial in $1/\varepsilon$. As the height of the semi-unfolding is at most exponential in the size of \mathcal{G} , the number of kernel approximations needed is at most doubly-exponential. The rest of the algorithm consist in applying Theorem 7.1, outside of the kernels. To obtain the overall complexity, we prove for kernel nodes an equivalent to Proposition 7.9, explaining how one step of computation on the semi-unfolding increases the complexity of the encoding of piecewise affine value functions.

Lemma 8.5. *For a kernel WTG $\tilde{\mathcal{G}}'$, if all piecewise affine value functions \mathbf{V}'_s of target children s' of s in T have complexity at most $\langle m, \beta \rangle$, then \mathbf{V}'_s , obtained by Lemma 8.3 with granularity N , has complexity at most $\langle m', \beta' \rangle$ with $m' = n(n+2)(MN+1)$ and $\beta' = (w_{\max} + n\beta)(n+1)!(MN)^{n+2}$, where n denotes the number of clocks.*

Proof. The bound on m' comes from the fact that value function \mathbf{V}'_s is obtained by interpolating the value computed for all $1/N$ -corners. The partition thus needs to separate each $1/N$ -region from others (in the worst case). By the characterisation following Definition 2.1, we thus need all (hyperplane) equations of the form $Nx_i = k$ with $k \in \{0, 1, \dots, MN\}$ (there are $n(MN+1)$ such equations), and $N(x_i - x_j) = k$ with $1 \leq i < j \leq n$ and $k \in \{-MN, -MN+1, \dots, MN\}$ (there are $n(n+1)/2 \times (2MN+1)$ such equations). In total, we thus need $n(MN+1) + n(n+1)(2MN+1)/2 \leq n(n+2)(MN+1)$ equations in the worst case.

Once multiplied by N , these equations use constants bounded by MN in absolute value. To bound β' , we also need to bound the constants appearing in the partition function F describing \mathbf{V}'_s on each base cell (that is a $1/N$ -region). We deduce from [BGHM16, Proposition 19] that the value of each of the corners in the $1/N$ -region game $\Gamma_N(\mathcal{G})$ is the value of an acyclic path in the untimed game $\Gamma_N(\mathcal{G})$ (since Max always has a memoryless optimal strategy), and is thus the sum of a final weight and of the weight of at most $(n+1)|L||\text{Reg}_N(\mathcal{X}, M)|$ transitions (this is the number of states of the untimed game).

- The final weight is obtained by taking the value of the partition function $F_{s'}$ of one of the children of s in T at a corner of a $1/N$ -region. If $F_{s'}$ is described by the equation $a_y y = a_1 x_1 + \dots + a_n x_n + b$ (with integers $|a_y|, |a_1|, \dots, |a_n|, |b| \leq \beta$) then the value at corner $(x_1 \mapsto k_1/N, \dots, x_n \mapsto k_n/N)$, with natural numbers $0 \leq k_1, \dots, k_n \leq NM$, is of the form $(a_1 k_1 + \dots + a_n k_n)/(Na_y) + b/a_y$. It can thus be written A/B with $|A|$ and $|B|$ integers at most $nMN\beta$.
- Each transition of $\Gamma_N(\mathcal{G})$ has a weight of the form $d \text{wt}(\ell) + \text{wt}(e)$ with d a delay that separates two corners: thus $d \in \{0, 1/N, 2/N, \dots, MN/N\}$, and the weight can be written as A/B with $|A|$ and $|B|$ integers at most $MNw_{\max}^L + Nw_{\max}^E$.

In total, corners have values that can be written as A/B with $|A|$ and $|B|$ integers at most $MN(w_{\max} + n\beta)$.

The partition function F describing \mathbf{V}'_s on a $1/N$ -region r is obtained by interpolating the values of each of its corners. Fix $\nu \in r$: it can be expressed as a (unique) convex combination of the $1/N$ -corners $\mathbf{v}_1, \dots, \mathbf{v}_{n+1}$ of r , so that $\nu = \sum_{i=1}^{n+1} \alpha_i \mathbf{v}_i$ with $\alpha_i \in [0, 1]$. Then, we let $\llbracket F \rrbracket(\nu) = \sum_{i=1}^{n+1} \alpha_i \text{Val}_{\Gamma_N(\mathcal{G})}((s, r, \mathbf{v}), \mathbf{v})$. To further bound β' , we need to express

the coefficients α_i in terms of $\nu(x_1), \dots, \nu(x_n)$. Indeed, we have the matricial relation

$$\begin{pmatrix} \nu(x_1) \\ \vdots \\ \nu(x_n) \\ 1 \end{pmatrix} = \begin{pmatrix} v_1(x_1) & v_2(x_1) & \cdots & v_{n+1}(x_1) \\ \vdots & \vdots & & \vdots \\ v_1(x_n) & v_2(x_n) & \cdots & v_{n+1}(x_n) \\ 1 & 1 & \cdots & 1 \end{pmatrix} \times \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_n \\ \alpha_{n+1} \end{pmatrix}$$

the last line coming from the fact that the sum of coefficients α_i must be equal to 1. The square matrix \mathcal{M} of this equality can be shown invertible, since no three corners of a $1/N$ -region are aligned. Therefore, each coefficient α_i can be written as an affine expression over $\nu(x_1), \dots, \nu(x_n)$ whose coefficients are given by the inverse of matrix \mathcal{M} . By Cramer's rule, this inverse can be written as the cofactor matrix \mathcal{C} divided by the determinant of \mathcal{M} . Coefficients of \mathcal{C} are determinants of $n \times n$ submatrices of \mathcal{M} . Since coefficients of \mathcal{M} are of the form k/N with $k \in \{0, \dots, MN\}$, coefficients of \mathcal{C} are of the form k/N^n with $k \in \{0, 1, \dots, n!(MN)^n\}$. The determinant of \mathcal{M} is similarly of the form k'/N^{n+1} with $k' \in \{1, \dots, (n+1)!(MN)^{n+1}\}$. We can thus write each α_i as an affine coefficient of $\nu(x_1), \dots, \nu(x_n)$ whose coefficients are of the form Nk/k' with $k \in \{0, 1, \dots, n!(MN)^n\}$ and $k' \in \{1, \dots, (n+1)!(MN)^{n+1}\}$. Multiplying these coefficients with the values of corners, we obtain that each partial derivative and additive constant of the equalities defining F on each $1/N$ -region is bounded by $MN(w_{\max} + n\beta)(n+1)!(MN)^{n+1}$. \square

Pairing this lemma with Proposition 7.9 that explains how the complexity grows along a non-kernel node, we obtain (as in Corollary 7.10) the maximal complexity of the partition functions obtained from the leaves of the semi-unfolding in i steps. Notice that the m -complexity is reset each time we go through a kernel, while the β -complexity continues to grow polynomially in β (but exponentially in n which does not change the final computations). Thus, we conclude once again that the partition functions after i steps have a complexity doubly-exponential in i , and thus (since the height h of the semi-unfolding is exponential in the size of the game), we obtain a triply-exponential algorithm computing an ε -approximation of the value of an almost-divergent WTG, which concludes the proof of Theorem 3.9. This complexity is polynomial in $1/\varepsilon$ as N is linear in $1/\varepsilon$. An example of execution of the approximation schema can be found in Appendix B.

9. SYMBOLIC ALGORITHMS

The previous approximation result suffers from several drawbacks. It relies on the SCC decomposition of the region automaton. Each of these SCCs have to be analysed in a sequential way, and their analysis requires an a priori refinement of the granularity of regions. We show that this can be overcome, in case we suppose that no configuration has value $-\infty$ (which could be guaranteed by the designer of the game for some particular reasons; this is for instance always the case if only non-negative weights are used). We prove in this section that the symbolic approach based on the value iteration paradigm, *i.e.* the computation of iterates of the operator \mathcal{F} , recalled in page 20, is an approximation schema, as stated in Theorem 3.10.

9.1. Symbolic approximation algorithm. Notice that configurations with value $+\infty$ are stable through value iteration, and do not affect its other computations. Since Theorem 3.10 assumes the absence of configurations of value $-\infty$, we will therefore consider in the following that all configurations have finite value in \mathcal{G} (we discuss further this hypothesis in Section 9.2).

Consider first a game \mathcal{G} that is a kernel, with final weight functions that are Λ -Lipschitz-continuous on regions. By Lemma 8.2, there exists an integer N such that solving the untimed weighted game $\Gamma_N(\mathcal{G})$ computes an $\varepsilon/2$ -approximation of the value of $1/N$ corners. We denote $N(\varepsilon, \Lambda)$ this granularity, and recall that

$$N(\varepsilon, \Lambda) = \left\lceil \frac{2w_{\max}^L |\mathcal{R}(\mathcal{G})| + 2\Lambda}{\varepsilon} \right\rceil.$$

Using the results of [BGHM16] for untimed weighted games, we know that those values are obtained after a finite number of steps of (the untimed version of) the value iteration operator, depending on the number $|\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})| \leq |L| |\text{Reg}_{N(\varepsilon, \Lambda)}(\mathcal{X}, M)| (n+1)$ of locations of the corner-point abstraction. More precisely, if one considers a number of iterations

$$P_{\mathcal{K}}(\varepsilon, \Lambda) = |\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})| (2(|\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})| - 1)w_{\max} + 1),$$

then $\text{Val}_{\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})}^{P_{\mathcal{K}}(\varepsilon, \Lambda)}((\ell, r, \nu), \nu) = \text{Val}_{\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})}((\ell, r, \nu), \nu)$. From this observation, we deduce the following property of $P_{\mathcal{K}}(\varepsilon, \Lambda)$:

Lemma 9.1. *If \mathcal{G} is a kernel with no configurations of infinite value and with final weight functions that are Λ -Lipschitz-continuous on regions, and $\varepsilon > 0$, then it holds for all configurations (ℓ, ν) of \mathcal{G} that $|\text{Val}_{\mathcal{G}}(\ell, \nu) - \text{Val}_{\mathcal{G}}^{P_{\mathcal{K}}(\varepsilon, \Lambda)}(\ell, \nu)| \leq \varepsilon$.*

Proof. We already know that $\text{Val}_{\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})}^{P_{\mathcal{K}}(\varepsilon, \Lambda)}((\ell, r, \nu), \nu) = \text{Val}_{\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})}((\ell, r, \nu), \nu)$ for all configurations $((\ell, r, \nu), \nu)$ of $\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})$. Moreover, Lemma 8.2 ensures that $|\text{Val}_{\mathcal{G}}(\ell, \nu) - \text{Val}_{\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})}((\ell, r, \nu), \nu)| \leq \varepsilon/2$ whenever ν is in the $\frac{1}{N(\varepsilon, \Lambda)}$ -region r . Therefore, we only need to prove that $|\text{Val}_{\mathcal{G}}^{P_{\mathcal{K}}(\varepsilon, \Lambda)}(\ell, \nu) - \text{Val}_{\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})}^{P_{\mathcal{K}}(\varepsilon, \Lambda)}((\ell, r, \nu), \nu)| \leq \varepsilon/2$ to conclude. This is done as for Lemma 8.2, since Lemma 8.1 (that we need to prove Lemma 8.2) does not depend on the length of the plays ρ and ρ' , and both runs reach the target state in the same step, *i.e.* both before or after the horizon of $P_{\mathcal{K}}(\varepsilon, \Lambda)$ steps. \square

Once we know that value iteration converges on kernels, we can use the semi-unfolding of Section 6 to prove that it also converges on non-negative SCCs when all values are finite.

We define the following notations:

- let h be the height of the semi-unfolding of \mathcal{G} ,
- let α be the maximum number of kernels in a branch of the semi-unfolding,
- let $(\Lambda_i)_{i \in \mathbb{N}}$ be a sequence of Lipschitz constants, defined by $\Lambda'_0 = \Lambda$ and by

$$\Lambda_{i+1} = \frac{\max(n-1, 2)^{P_{\mathcal{K}}(\frac{\varepsilon}{\alpha}, \Lambda_i)} - 1}{\max(n-2, 1)} w_{\max}^L + (n-1)^{P_{\mathcal{K}}(\frac{\varepsilon}{\alpha}, \Lambda_i)} \Lambda_i,$$

- a number of iterations $P_+(\varepsilon, \Lambda) = h + \alpha(P_{\mathcal{K}}(\frac{\varepsilon}{\alpha}, \Lambda_h) - 1)$.

In the worst case, $P_+(\varepsilon, \Lambda)$ can be non-elementary: $P_{\mathcal{K}}(\frac{\varepsilon}{\alpha}, \Lambda)$ grows polynomially in Λ , so that Λ_h is a tower of α exponentials.

Lemma 9.2. *If \mathcal{G} is a non-negative SCC with no configurations of infinite value and with final weight functions that are Λ -Lipschitz-continuous on regions, and $\varepsilon > 0$, then $|\text{Val}_{\mathcal{G}}(\ell, \nu) - \text{Val}_{\mathcal{G}}^{P_+(\varepsilon, \Lambda)}(\ell, \nu)| \leq \varepsilon$ for all configurations (ℓ, ν) of \mathcal{G} .*

Proof. The idea is to unfold every kernel of the semi-unfolding game $\mathcal{T}(\mathcal{G})$ according to its bound in Lemma 9.1. More precisely, consider a non-negative SCC \mathcal{G} , a precision ε , and an initial configuration (ℓ_0, ν_0) . Let $\mathcal{T}(\mathcal{G})$ be its finite semi-unfolding (obtained from the labelled tree T , as in Section 6), such that $\text{Val}_{\mathcal{G}}(\ell_0, \nu_0) = \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$. Let α be the maximum number of kernels along one of the branches of T .

We describe a bottom-up transformation of T into T' , that keeps track of the Lipschitz constants and the precision of the approximation, and define a new weighted timed game $\mathcal{T}'(\mathcal{G})$ from T' by applying the method used to create $\mathcal{T}(\mathcal{G})$ in Section 6.

Leaf nodes are left unchanged, with final weight functions that are Λ -Lipschitz-continuous on regions. Note that the associated state in $\mathcal{T}(\mathcal{G})$ has the same value. For an internal node s of T (that gives rise to a single state (ℓ, r) of $\mathcal{T}(\mathcal{G})$, or a kernel $\mathsf{K}_{\ell, r}$), let us suppose that the value of all children s' of s in T' are Λ_i -Lipschitz-continuous on regions, for some Λ_i , and that they are ε_i -close to their values in $\mathcal{T}(\mathcal{G})$.

If s is a node of the form (ℓ, r) (that is not part of a kernel), we keep it unchanged in T' , and therefore by Lemma 8.4 the value of s in $\mathcal{T}'(\mathcal{G})$ is ε_i -close to its value in $\mathcal{T}(\mathcal{G})$, and is $\max(\Lambda_i, |\text{wt}(\ell)| + (n-1)\Lambda_i)$ -Lipschitz-continuous on regions by Lemma 7.13.

If s is a node of the form $\mathsf{K}_{\ell, r}$, we replace it by an unfolding of $\mathsf{K}_{\ell, r}$ of depth $P_{\mathsf{K}}(\frac{\varepsilon}{\alpha}, \Lambda_i)$, where P is the bound of Lemma 9.1 for the kernel $\mathsf{K}_{\ell, r}$. Then, by Lemmas 9.1 and 8.4, the value of s in $\mathcal{T}'(\mathcal{G})$ is $\varepsilon_i + \frac{\varepsilon}{\alpha}$ -close to its value in $\mathcal{T}(\mathcal{G})$, and is Λ_{i+1} -Lipschitz-continuous on regions by Proposition 7.12, with

$$\Lambda_{i+1} = \frac{\max(n-1, 2)^{P_{\mathsf{K}}(\frac{\varepsilon}{\alpha}, \Lambda_i)} - 1}{\max(n-2, 1)} w_{\max}^L + (n-1)^{P_{\mathsf{K}}(\frac{\varepsilon}{\alpha}, \Lambda_i)} \Lambda_i.$$

In particular, note that $P_{\mathsf{K}}(\frac{\varepsilon}{\alpha}, \Lambda_i)$ is polynomial in Λ_i , therefore Λ_{i+1} is exponential in Λ_i . Since $P_{\mathsf{K}}(\frac{\varepsilon}{\alpha}, \Lambda_i) \geq 1$, Λ_{i+1} is greater than the $\max(\Lambda_i, |\text{wt}(\ell)| + (n-1)\Lambda_i)$ expression obtained for non-kernel nodes. Thus, we can bound the Lipschitz constants globally by Λ_h , so that each kernel is unfolded for at most $P_{\mathsf{K}}(\frac{\varepsilon}{\alpha}, \Lambda_h)$ steps.

At the root of T' , we recover a node whose value function in $\mathcal{T}'(\mathcal{G})$ is ε -close to its value in $\mathcal{T}(\mathcal{G})$, and that is Λ_h -Lipschitz-continuous on regions.

Observe that $\mathcal{T}'(\mathcal{G})$ is not a semi-unfolding, it is instead a (complete) unfolding of $\mathcal{R}(\mathcal{G})$, of a certain bounded depth $P_+(\varepsilon, \Lambda)$ (at most $(h-\alpha)$ plus α times $P_{\mathsf{K}}(\frac{\varepsilon}{\alpha}, \Lambda_h)$).

Finally, let us show that the value computed by $\text{Val}_{\mathcal{G}}^{P_+(\varepsilon, \Lambda)}$ on the root state $(\tilde{\ell}_0, r_0)$ is bounded between $\text{Val}_{\mathcal{G}}$ and $\text{Val}_{\mathcal{T}'(\mathcal{G})}$, which allows us to conclude. Consider $\mathcal{T}''(\mathcal{G})$, the (complete) unfolding of $\mathcal{R}(\mathcal{G})$ with unfolding depth $P_+(\varepsilon, \Lambda)$, where kernels are also unfolded. By construction, $\text{Val}_{\mathcal{T}''(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) = \text{Val}_{\mathcal{T}''(\mathcal{G})}^{P_+(\varepsilon, \Lambda)}((\tilde{\ell}_0, r_0), \nu_0)$. Then, we can prove that $\text{Val}_{\mathcal{T}''(\mathcal{G})}^{P_+(\varepsilon, \Lambda)}((\tilde{\ell}_0, r_0), \nu_0) = \text{Val}_{\mathcal{G}}^{P_+(\varepsilon, \Lambda)}(\ell_0, \nu_0)$ (same strategies at bounded horizon $P_+(\varepsilon, \Lambda)$), which implies $\text{Val}_{\mathcal{G}}(\ell_0, \nu_0) = \text{Val}_{\mathcal{R}(\mathcal{G})}((\ell_0, r_0), \nu_0) \leq \text{Val}_{\mathcal{T}''(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$ (the function Val^k decreases monotonously as k increases). By another monotonicity argument, we can also prove $\text{Val}_{\mathcal{T}''(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) \leq \text{Val}_{\mathcal{T}'(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$: the tree $\mathcal{T}''(\mathcal{G})$ contains $\mathcal{T}'(\mathcal{G})$ as a prefix, and therefore one can construct $\mathcal{T}'(\mathcal{G})$ from $\mathcal{T}''(\mathcal{G})$ by replacing sub-trees of $\mathcal{T}''(\mathcal{G})$ by stopped leaves of final weight $+\infty$ (remember that we are in a non-negative SCCs). Bringing everything together we obtain $|\text{Val}_{\mathcal{G}}^{P_+(\varepsilon, \Lambda)}(\ell_0, \nu_0) - \text{Val}_{\mathcal{G}}(\ell_0, \nu_0)| \leq \varepsilon$. \square

Proving the same property on non-positive SCCs requires more work, because the semi-unfolding gives final weight $-\infty$ to stop leaves, which does not integrate well with value iteration (initialisation at $+\infty$ on non-target states). However, by unfolding those SCCs slightly more (at most $|\mathcal{R}(\mathcal{G})|$ more steps), we can obtain the desired property with a similar bound $P_-(\varepsilon, \Lambda) = h + |\mathcal{R}(\mathcal{G})| + \alpha(P_K(\frac{\varepsilon}{\alpha}, \Lambda_{h+|\mathcal{R}(\mathcal{G})|}) - 1)$.

Lemma 9.3. *If \mathcal{G} is a non-positive SCC with no configurations of infinite value and with final weight functions that are Λ -Lipschitz-continuous on regions, and $\varepsilon > 0$, then $|\text{Val}_{\mathcal{G}}(\ell, \nu) - \text{Val}_{\mathcal{G}}^{P_-(\varepsilon, \Lambda)}(\ell, \nu)| \leq \varepsilon$ for all configurations (ℓ, ν) of \mathcal{G} .*

Proof. Consider a non-positive SCC \mathcal{G} , a precision ε , and an initial configuration (ℓ_0, ν_0) . Let $\mathcal{T}(\mathcal{G})$ be its finite semi-unfolding (obtained from the labelled tree T , as in Section 6), such that $\text{Val}_{\mathcal{G}}(\ell_0, \nu_0) = \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$.

We now change T , by adding a subtree under each stopped leaf: the complete unfolding of $\mathcal{R}(\mathcal{G})$, starting from the stopped leaf, of depth $|\mathcal{R}(\mathcal{G})|$. Let us name T^+ this unfolding tree. We then construct $\mathcal{T}^+(\mathcal{G})$ as before, based on T^+ . Since we are in a non-positive SCC, $\mathcal{T}^+(\mathcal{G})$ must have final weight $-\infty$ on its stopped leaves. It is easy to see that $\text{Val}_{\mathcal{G}}(\ell_0, \nu_0) = \text{Val}_{\mathcal{T}^+(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$ still holds (the proof of Lemma 9.2 was based on branches being long enough, and we increased the lengths). We now perform a small but crucial change: the final weight of stopped leaves in $\mathcal{T}^+(\mathcal{G})$ is set to $+\infty$ instead of $-\infty$. Trivially $\text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) \leq \text{Val}_{\mathcal{T}^+(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$ (we increased the final weight function). Let us prove that

$$\text{Val}_{\mathcal{T}^+(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) \leq \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0).$$

For a fixed $\eta > 0$, consider σ_{Min} an η -optimal strategy for player Min in $\mathcal{T}(\mathcal{G})$. Let us define σ_{Min}^+ , a strategy for Min in $\mathcal{T}^+(\mathcal{G})$, by making the same choice as σ_{Min} on the common prefix tree, and once a node that is a stopped leaf in $\mathcal{T}(\mathcal{G})$ is reached, we switch to a positional attractor strategy of Min towards target states. Consider any strategy σ_{Max}^+ of Max in $\mathcal{T}^+(\mathcal{G})$, and let σ_{Max} be its projection in $\mathcal{T}(\mathcal{G})$. Let ρ^+ denote the (maximal) play

$$\text{play}_{\mathcal{T}^+(\mathcal{G})}(((\ell_0, r_0), \nu_0), \sigma_{\text{Min}}^+, \sigma_{\text{Max}}^+),$$

and ρ be $\text{play}_{\mathcal{T}(\mathcal{G})}(((\ell_0, r_0), \nu_0), \sigma_{\text{Min}}, \sigma_{\text{Max}})$. By construction, ρ^+ does not reach a stopped leaf in $\mathcal{T}^+(\mathcal{G})$. If the play ρ^+ stays in the common prefix tree of T and T^+ , then $\rho = \rho^+$, and

$$\text{wt}_{\mathcal{T}^+(\mathcal{G})}(\rho^+) \leq \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) + \eta.$$

If it does not, then ρ^+ has a prefix that reaches a stopped leaf in $\mathcal{T}(\mathcal{G})$: this must be ρ . This implies (see Lemma 6.4) that

$$\text{wt}_{\mathcal{T}^+(\mathcal{G})}(\rho^+) < -|\mathcal{R}(\mathcal{G})|w_{\text{max}} - w_{\text{max}}^{\dagger} \leq \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0).$$

Since this holds for all $\eta > 0$, we proved $\text{Val}_{\mathcal{T}^+(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0) \leq \text{Val}_{\mathcal{T}(\mathcal{G})}((\tilde{\ell}_0, r_0), \nu_0)$, which finally implies that the two values are equal.

Then, we can follow the proof of Lemma 9.2 (with T^+ and $\mathcal{T}^+(\mathcal{G})$ instead of T and $\mathcal{T}(\mathcal{G})$) in order to conclude. \square

Now, if we are given an almost-divergent WTG \mathcal{G} and a precision ε , we can add the bounds for value iteration obtained from each SCC by Lemmas 9.2 and 9.3, and obtain a final bound P such that for all $k \geq P$, $\text{Val}_{\mathcal{G}}^k$ is an ε -approximation of $\text{Val}_{\mathcal{G}}$. This concludes the proof of Theorem 3.10.

Remark 9.4. We note that one can get a (non-elementary) upper bound for P without constructing the semi-unfolding nor the region game, by using the inequalities $|\mathcal{R}(\mathcal{G})| \leq |L|n!(2M)^n$, $|\Gamma_{N(\varepsilon, \Lambda)}(\mathcal{G})| \leq |L|n!(2N(\varepsilon, \Lambda)M)^n(n+1)$ and $\alpha \leq h \leq |\mathcal{R}(\mathcal{G})|(3|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 2) + 1$ to fix global values for $\alpha = h$, followed by observing $P_+(\varepsilon, \Lambda) \leq P_-(\varepsilon, \Lambda)$ so that $P \leq |\mathcal{R}(\mathcal{G})|P_-(\varepsilon, \Lambda)$.

9.2. Discussion. This symbolic procedure avoids the three drawbacks (SCC decomposition, sequential analysis of the SCCs, and refinement of the granularity of regions) of the previous approximation schema if we assume that no configuration has value $-\infty$. However, computing the number of steps P of Theorem 3.10 is non-elementary, with an upper bound complexity that is of the order of a tower of α exponentials, with α exponential in the size of the input WTG. Instead, we could directly launch the value iteration algorithm on the game \mathcal{G} , and stop the computation whenever we are satisfied enough by the approximation computed: however, there are no formal guarantees whatsoever on the quality of the approximation before the number of steps P given above.

If \mathcal{G} is not guaranteed to be free of configurations of value $-\infty$, then we must first perform the SCC decomposition of $\mathcal{R}(\mathcal{G})$, and, as \mathcal{G} is almost-divergent, identify and remove regions whose value is $-\infty$, by Proposition 5.5. Then, we can apply the value iteration algorithm.

We also note that if \mathcal{G} is a divergent WTG, the unfolding of kernels is not needed, so that Lemmas 9.2 and 9.3 construct bounds $P_+(\varepsilon, \Lambda)$ and $P_-(\varepsilon, \Lambda)$ allowing one to compute the exact values of \mathcal{G} . Moreover, these bounds become exponential in the size of \mathcal{G} instead of being non-elementary, so that the overall complexity of the symbolic algorithm is 3-EXPTIME, matching the result of Section 8.

As a final remark, notice that our correctness proof strongly relies on Section 8.2, and thus would not hold with the approximation schema of [BJM15] (which does not preserve the continuity on regions of the computed value functions, in turn needed to define final weights on $\frac{1}{N(\varepsilon, \Lambda)}$ -corners).

10. STRATEGY SYNTHESIS

We are also interested in the synthesis problem, that asks for an ε -optimal strategy of Min . In this section, we will prove the following result:

Theorem 10.1. *Let $\varepsilon > 0$ and let \mathcal{G} be a divergent WTG. We can compute in triple exponential time an ε -optimal strategy for Min .*

Recall that in the value iteration algorithm of Section 2.8, one step of the game is summarised by the operator \mathcal{F} mapping each value function \mathbf{V} to a value function $\mathbf{V}' = \mathcal{F}(\mathbf{V})$ defined by

$$\mathbf{V}'_{\ell}(\nu) = \begin{cases} \text{fwt}(\ell, \nu) & \text{if } \ell \in L_{\text{t}}, \\ \sup_{(\ell, \nu) \xrightarrow{d, e} (\ell', \nu')} [d \cdot \text{wt}(\ell) + \text{wt}(e) + \mathbf{V}_{\ell'}(\nu')] & \text{if } \ell \in L_{\text{Max}}, \\ \inf_{(\ell, \nu) \xrightarrow{d, e} (\ell', \nu')} [d \cdot \text{wt}(\ell) + \text{wt}(e) + \mathbf{V}_{\ell'}(\nu')] & \text{if } \ell \in L_{\text{Min}}. \end{cases} \quad (10.1)$$

Intuitively, ε -optimal strategies can be extracted from the value iteration operator, by mapping each play that ends in (ℓ, ν) with $\ell \in L_{\text{Min}}$ to a choice (d, e) such that the transition

$(\ell, \nu) \xrightarrow{d,e} (\ell', \nu')$ is ε -optimal. However, the choice depends on the step i in the value iteration computation. Formally, if A is a set, f is a mapping from A to \mathbb{R}_∞ and $\varepsilon > 0$, let $\operatorname{arginf}_{a \in A}^\varepsilon f(a)$ denote the set of elements $a^* \in A$ such that $f(a^*) \leq \inf_{a \in A} f(a) + \varepsilon$. Then, let us name $\sigma_{\text{Min}}^{i,\varepsilon}$ a strategy defined from the application of (10.1) in $\mathbf{V}^i = \mathcal{F}(\mathbf{V}^{i-1})$, so that for all finite plays ρ ending in (ℓ, ν) and $\ell \in L_{\text{Min}}$,

$$\sigma_{\text{Min}}^{i,\varepsilon}(\rho) \in \operatorname{arginf}_{(\ell, \nu) \xrightarrow{d,e} (\ell', \nu')}^\varepsilon (d \cdot \operatorname{wt}(\ell) + \operatorname{wt}(e) + \mathbf{V}_{\ell'}^{i-1}(\nu')) \quad (10.2)$$

Now, let $\sigma_{\text{Min}}^{*,i,\varepsilon}$ denote a strategy that maps every finite play ρ ending in $L_{\text{Min}} \setminus L_{\text{t}}$ to $(d, e) = \sigma_{\text{Min}}^{j,\varepsilon}(\rho)$, with $j = \max(0, i - |\rho|)$.

Proposition 10.2. *The strategy $\sigma_{\text{Min}}^{*,i,\varepsilon}$ is εi -optimal for Min during the first i steps:*

$$|\operatorname{Val}^i((\ell, \nu), \sigma_{\text{Min}}^{*,i,\varepsilon/i}) - \operatorname{Val}^i(\ell, \nu)| \leq \varepsilon$$

Proof. Let us show by induction on i that for all configurations (ℓ, ν) in \mathcal{G} , $\operatorname{Val}_{\mathcal{G}}^i((\ell, \nu), \sigma_{\text{Min}}^{*,i,\varepsilon}) \leq \mathbf{V}_{\ell}^i(\nu) + \varepsilon i$. If $i = 0$, both sides are equal to either $+\infty$ or to $\operatorname{fwt}(\ell, \nu)$ if $\ell \in L_{\text{t}}$. Let us assume that $\operatorname{Val}_{\mathcal{G}}^{i-1}((\ell, \nu), \sigma_{\text{Min}}^{*,i-1,\varepsilon}) \leq \mathbf{V}_{\ell}^{i-1}(\nu) + \varepsilon(i-1)$.

We note that for all strategies σ_{Min} , if $\ell \in L_{\text{Max}}$ then

$$\operatorname{Val}_{\mathcal{G}}^i((\ell, \nu), \sigma_{\text{Min}}) = \sup_{(\ell, \nu) \xrightarrow{d,e} (\ell', \nu')} [d \cdot \operatorname{wt}(\ell) + \operatorname{wt}(e) + \operatorname{Val}_{\mathcal{G}}^{i-1}((\ell', \nu'), \sigma'_{\text{Min}})],$$

where σ'_{Min} appends $(\ell, \nu) \xrightarrow{d,e} (\ell', \nu')$ in front of paths and then calls σ_{Min} . In particular, if $\sigma_{\text{Min}} = \sigma_{\text{Min}}^{*,i,\varepsilon}$ then σ'_{Min} matches the definition of $\sigma_{\text{Min}}^{*,i-1,\varepsilon}$. Then, $\operatorname{Val}_{\mathcal{G}}^i((\ell, \nu), \sigma_{\text{Min}}^{*,i,\varepsilon}) \leq \mathbf{V}_{\ell}^i(\nu) + \varepsilon(i-1)$.

Similarly, if $\ell \in L_{\text{Min}} \setminus L_{\text{t}}$ then

$$\operatorname{Val}_{\mathcal{G}}^i((\ell, \nu), \sigma_{\text{Min}}) = d \cdot \operatorname{wt}(\ell) + \operatorname{wt}(e) + \operatorname{Val}_{\mathcal{G}}^{i-1}((\ell', \nu'), \sigma'_{\text{Min}})$$

where $(\ell, \nu) \xrightarrow{d,e} (\ell', \nu')$ is the decision made by σ_{Min} on (ℓ, ν) and σ'_{Min} appends $(\ell, \nu) \xrightarrow{d,e} (\ell', \nu')$ in front of paths and then calls σ_{Min} . Then,

$$\operatorname{Val}_{\mathcal{G}}^i((\ell, \nu), \sigma_{\text{Min}}^{*,i,\varepsilon}) - \varepsilon \leq d \cdot \operatorname{wt}(\ell) + \operatorname{wt}(e) + \operatorname{Val}_{\mathcal{G}}^{i-1}((\ell', \nu'), \sigma_{\text{Min}}^{*,i-1,\varepsilon}) \leq \mathbf{V}_{\ell}^i(\nu) + \varepsilon(i-1)$$

Finally, if $\ell \in L_{\text{t}}$ then $\operatorname{Val}_{\mathcal{G}}^i((\ell, \nu), \sigma_{\text{Min}}^{*,i,\varepsilon}) = \mathbf{V}_{\ell}^i(\nu) = \operatorname{fwt}(\ell, \nu)$. \square

We now explain how to extract some strategies $\sigma_{\text{Min}}^{i,\varepsilon}$ from the partition functions, in order to solve the synthesis problems on acyclic and divergent games (games where a known $i \in \mathbb{N}$ implies $\operatorname{Val} = \operatorname{Val}^i$).

We will use affine inequalities over $n+1$ variables to encode constraints on the choice of delays. Formally, an affine delay inequality is an equation I of the form $a_{\text{d}}\mathbf{d} < a_1\mathbf{x}_1 + \dots + a_n\mathbf{x}_n + b$, with all a_i, b and a_{d} integers of \mathbb{Z} , $a_{\text{d}} \neq 0$ and $< \in \{<, \leq\}$. We say that I is a lower bound if $a_{\text{d}} < 0$, and that it is an upper bound if $a_{\text{d}} > 0$.

Definition 10.3. A *partition strategy function* S defined over a partition P is a mapping from the base cells of P to a tuple $(e, I_{\mathbf{l}}, I_{\mathbf{u}}, p)$, where e is an edge of \mathcal{G} , $I_{\mathbf{l}}$ and $I_{\mathbf{u}}$ are two affine delay inequalities that are respectively lower and upper bounds, and finally $p \in \{\mathbf{l}, \mathbf{u}\}$ selects one of the two.

Given a precision $\varepsilon > 0$, a partition strategy function encodes a mapping from c_P to pairs (e, d^ε) denoted $\llbracket S \rrbracket^\varepsilon$, with $e \in E$ and $d^\varepsilon \in \mathbb{R}_{\geq 0}$. Let $\nu \in c_P$ and $S([\nu]_P)$ be equal to (e, I_1, I_u, p) , with I_1 defined by $a_d \mathbf{d} \prec a_1 x_1 + \dots + a_n x_n + b$ and I_u defined by $a'_d \mathbf{d} \prec' a'_1 x_1 + \dots + a'_n x_n + b'$. Let $D \subseteq \mathbb{R}_{\geq 0}$ denote the interval of delays \mathbf{d} that satisfy both $a_d \mathbf{d} \prec a_1 \nu(x_1) + \dots + a_n \nu(x_n) + b$ and $a'_d \mathbf{d} \prec' a'_1 \nu(x_1) + \dots + a'_n \nu(x_n) + b'$. We denote by d_1 and d_u the lower and upper bounds of D . Let D^ε denote $D \cap (d_1 - \varepsilon, d_1 + \varepsilon)$ if $p = \mathbf{1}$, and $D \cap (d_u - \varepsilon, d_u + \varepsilon)$ if $p = \mathbf{u}$. We denote by d_1^ε and d_u^ε the lower and upper bounds of D^ε . Then, $\llbracket S \rrbracket^\varepsilon(\nu) = (e, d^\varepsilon)$ with $d^\varepsilon = \frac{d_1^\varepsilon + d_u^\varepsilon}{2} \in D^\varepsilon$.

We argue that partition strategy functions can be used to compute a positional strategy $\sigma_{\text{Min}}^{i,\varepsilon}$ given an encoding of $\text{Val}_{\mathcal{G}}^{i-1}$ as partition functions, so that (10.2) holds.

We assume that for all $\ell \in L$, \mathbf{V}_ℓ^i is piecewise affine with finitely many pieces and is Λ -Lipschitz-continuous over regions. By Proposition 7.12, we know that Λ is at most exponential in the size of \mathcal{G} and i .

Proposition 10.4. *For all $\varepsilon > 0$ and $i \in \mathbb{N}_{>0}$, we can compute in time doubly-exponential in i and exponential in the size of \mathcal{G} a partition P_ℓ^i and a partition strategy function S_ℓ^i for each location $\ell \in L_{\text{Min}} \setminus L_t$, so that $\llbracket S_\ell^i \rrbracket^{\varepsilon, \Lambda}(\nu) = \sigma_{\text{Min}}^{i,\varepsilon}(\ell, \nu)$.*

Proof. In Section 7, we detailed how to compute piecewise affine value functions (P_ℓ^i, F_ℓ^i) encoding \mathbf{V}_ℓ^i for all locations ℓ and a fixed horizon i . We will now explain how S_ℓ^i can be obtained as a side-product of the step $\mathbf{V}^i = \mathcal{F}(\mathbf{V}^{i-1})$.

Recall that if $\ell \in L_{\text{Min}} \setminus L_t$, it holds that

$$\mathbf{V}_\ell^i = \min_{e=(\ell, g, \mathcal{Y}, \ell')} \text{Pre}_e(\text{Guard}_g(\text{Unreset}_{\mathcal{Y}}(\mathbf{V}_{\ell'}^{i-1})))$$

where e ranges over the edges in \mathcal{G} that starts from ℓ . We can identify, for every base cell c of the resulting partition P_ℓ^i , an edge e so that \mathbf{V}_ℓ^i restricted to domain c is equal to $\text{Pre}_e(\text{Guard}_g(\text{Unreset}_{\mathcal{Y}}(\mathbf{V}_{\ell'}^{i-1})))$. This is the edge that S_ℓ^i will play on c . For the choice of delay, let us explain how to obtain the delay inequalities I_1 and I_u , and the selection choice p , from the computation of Pre_e .

Let (P, F) be an atomic tube partition encoding $\text{Guard}_g(\text{Unreset}_{\mathcal{Y}}(\mathbf{V}_{\ell'}^{i-1}))$. If c_b is a base cell of P , recall that \mathcal{C}_{c_b} is the set of cells reachable from c_b by time elapse, and that if $c \in \mathcal{C}_{c_b}$, $\mathcal{B}_{c_b}(c)$ is the set of borders of c reachable from c_b by time-elapse. In particular, $\mathcal{B}_{c_b}(c)$ contains either one non-diagonal border (a single choice of delay reaches c from any $\nu \in c_b$) or two non-diagonal borders (in which case there is a range of delays reaching c). Then, $\text{Pre}_e(\llbracket F \rrbracket)$ restricted to domain c_b equals $\min(\llbracket F \rrbracket, \min_{c \in \mathcal{C}_{c_b}} \min_{B \in \mathcal{B}_{c_b}(c)} [\text{Pre}_{e,c,B}(F)])$. These minimum operators are applied over piecewise affine value functions, and we can identify, for every base cell c' of the output, the arguments that optimise them. In particular, the base cells c' where $\text{Pre}_e(\llbracket F \rrbracket) = \llbracket F \rrbracket$ correspond to a choice of delay zero, so that S_ℓ^i maps c' to $(e, -d \leq 0, \mathbf{d} \leq 0, \mathbf{1})$. Alternatively, the base cells c' where $\text{Pre}_e(\llbracket F \rrbracket) = \text{Pre}_{e,c,B}(F)$ for some $c \in \mathcal{C}_{c_b}$ and some $B \in \mathcal{B}_{c_b}(c)$ correspond, for all $\nu \in c'$, to a choice of delay arbitrarily close to $d_{\nu,B}$, the delay that lets valuation ν reach $\llbracket B \rrbracket$ by time-elapse. Recall that if B is described by $a_1 x_1 + \dots + a_n x_n + b = 0$ and $A = a_1 + \dots + a_n$, then

$$-Ad_{\nu,B} = a_1 \cdot \nu(x_1) + \dots + a_n \cdot \nu(x_n) + b$$

There are now two cases. If $\mathcal{B}_{c_b}(c) = \{B\}$, then S_ℓ^i maps c' to $(e, I_1, I_u, \mathbf{1})$, with I_1 defined by $-Ad \leq a_1 x_1 + \dots + a_n x_n + b$ and I_u defined by $Ad \leq -a_1 x_1 - \dots - a_n x_n - b$. Otherwise $\mathcal{B}_{c_b}(c) =$

$\{B_{\mathbf{l}}, B_{\mathbf{u}}\}$, with $B_{\mathbf{l}}$ the lower border and $B_{\mathbf{u}}$ the upper one.¹⁹ As $B_{\mathbf{l}}$ and $B_{\mathbf{u}}$ are borders of c , they correspond to affine inequalities $a_1x_1 + \dots + a_nx_n + b < 0$ and $a'_1x_1 + \dots + a'_nx_n + b' < 0$, respectively. Let $A = a_1 + \dots + a_n$ and $A' = a'_1 + \dots + a'_n$. We let $I_{\mathbf{l}}$ be defined by $-Ad < a_1x_1 + \dots + a_nx_n + b$ and $I_{\mathbf{u}}$ be defined by $A'd < -a'_1x_1 - \dots - a'_nx_n - b'$. Then we let S_{ℓ}^i map c' to $(e, I_{\mathbf{l}}, I_{\mathbf{u}}, p)$, where $p = \mathbf{l}$ if $B = B_{\mathbf{l}}$ and $p = \mathbf{u}$ if $B = B_{\mathbf{u}}$.

It follows that $\llbracket S_{\ell}^i \rrbracket^{\varepsilon}(\nu)$ corresponds to a choice of edge e and delay d from (ℓ, ν) leading to a valuation (ℓ', ν') , so that $\|(\nu + d) - (\nu + d_{\nu, B})\|_{\infty} < \varepsilon$, with $d_{\nu, B}$ the optimal delay for minimizing the Pre_e step. As the output piecewise affine functions are Λ -Lipschitz-continuous over regions, $|\mathbf{V}_{\ell}^i(\nu + d) - \mathbf{V}_{\ell'}^i(\nu + d_{\nu, B})| < \Lambda\varepsilon$. Then by definition of Pre_e we obtain

$$\llbracket S_{\ell}^i \rrbracket^{\varepsilon}(\nu) \in \underset{(\ell, \nu) \xrightarrow{d, e} (\ell', \nu')}{\text{arginf}}^{\Lambda\varepsilon} (d \cdot \text{wt}(\ell) + \text{wt}(e) + \mathbf{V}_{\ell'}^{i-1}(\nu')) \quad \square$$

By Proposition 10.2, we can therefore solve the synthesis problem in triple exponential time for all weighted time games \mathcal{G} so that $\text{Val}_{\mathcal{G}} = \text{Val}_{\mathcal{G}}^i$ with i at most exponential in the size of \mathcal{G} . This holds for acyclic games, where i is bounded by $|\mathcal{R}(\mathcal{G})|$, and also for divergent games, where i can be bounded by the results of Section 9.1 (Lemmas 9.2 and 9.3 in the special case with no kernel), assuming that an exponential time pre-computation using Proposition 5.5 is used to remove the valuations of value $-\infty$. This concludes the proof of Theorem 10.1.

11. WEIGHTED TIMED GAMES WITH NO CLOCKS

In this section, we study weighted timed games where the set of clocks is empty. Their semantics is a finite transition system, so we will simply call them (finite) weighted games. For notational convenience, we omit guards, resets, valuations and delays from all notations, and merge the notions of locations and configurations into the notion of states. We also assume that all delays are null and that weights are integers.

Definition 11.1. A *weighted game*²⁰ is a tuple $\mathcal{G} = \langle S_{\text{Min}}, S_{\text{Max}}, S_{\text{t}}, T, \text{wt}, \text{fwt} \rangle$ with $S = S_{\text{Min}} \uplus S_{\text{Max}}$ a finite set of states split between players Min and Max , $T \subseteq S \times S$ a finite set of transitions $s \rightarrow s'$ from state s to state s' , $\text{wt}: T \rightarrow \mathbb{Z}$ a weight function associating an integer weight with each transition²¹, $S_{\text{t}} \subseteq S_{\text{Min}}$ a set of target states for player Min , and $\text{fwt}: S_{\text{t}} \rightarrow \mathbb{Z}_{\infty}$ is a function mapping each target state to a final weight of $\mathbb{Z}_{\infty} = \mathbb{Z} \cup \{-\infty, +\infty\}$.

A cycle is a finite play ρ , of length at least 1, such that $\text{first}(\rho) = \text{last}(\rho)$.

11.1. Solving weighted games. Let \mathcal{G} be a finite weighted game. The value iteration algorithm described in Section 2.8 can be applied on \mathcal{G} as a greatest fixpoint computation over a vector of $|S|$ numbers, as detailed in [BGHM16]. This algorithm computes the value of any finite weighted game in pseudo-polynomial time. In particular, along the fixpoint computation states with a value of $-\infty$ will be associated with a sequence of values that converges towards $-\infty$. When the value of such a state gets low enough, it is recognised as a $-\infty$ state and is set to its fixpoint value directly.

¹⁹We can assume without loss of generality that $d_{\nu, B_{\mathbf{l}}} \leq d_{\nu, B_{\mathbf{u}}}$ for all $\nu \in c_b$ as the tube partition (P, F) is atomic.

²⁰Weighted games are called *min-cost reachability games* in [BGHM16].

²¹If $t = s \rightarrow s'$ is a transition in \mathcal{G} , we denote $\text{wt}(s, s')$ the weight $\text{wt}(t)$.

11.1.1. *Optimal strategies.* Let us fix an initial state s . By definition of lower and upper values, there exists for each player $P \in \{\text{Min}, \text{Max}\}$ a sequence of strategies $(\sigma_P^i)_{i \in \mathbb{N}}$ such that $\lim_{i \rightarrow \infty} \text{Val}_{\mathcal{G}}(s, \sigma_P^i) = \text{Val}_{\mathcal{G}}(s)$ and such that the sequence $(\text{Val}_{\mathcal{G}}(s, \sigma_{\text{Min}}^i))_{i \in \mathbb{N}}$ is non-increasing over \mathbb{Z}_{∞} , while $(\text{Val}_{\mathcal{G}}(s, \sigma_{\text{Max}}^i))_{i \in \mathbb{N}}$ is non-decreasing. If the sequence $(\text{Val}_{\mathcal{G}}(s, \sigma_P^i))_{i \in \mathbb{N}}$ stabilizes for all $i \geq k$, then σ_P^k is an optimal strategy of player P for s , *i.e.* $\text{Val}_{\mathcal{G}}(s, \sigma_P^k) = \text{Val}_{\mathcal{G}}(s)$. Therefore, if $\text{Val}_{\mathcal{G}}(s) > -\infty$ then Min must have an optimal strategy for s (an infinite decreasing sequence over \mathbb{Z} stabilizes), and if $\text{Val}_{\mathcal{G}}(s) < +\infty$ then Max has an optimal strategy for s . Moreover, if an optimal strategy for P exists for all states s , then they can be combined into an (overall) optimal strategy for P .

In fact, there always exists a *positional* strategy σ_{Max}^* for Max that is optimal, even if some states have value $+\infty$. The strategy σ_{Max}^* can be obtained in the value iteration algorithm, by memorizing for every state $s \in S_{\text{Max}}$ the transition that maximizes $\max_{s \rightarrow s'} [\text{wt}(s, s') + V_{s'}]$ in the last application of \mathcal{F} . However, this does not hold for Min , as there might be no sequence of positional strategies for player Min whose value at s converges towards $\text{Val}_{\mathcal{G}}(s)$. In [BGHM16], it is shown that value iteration can also compute an optimal strategy for Min (or a sequence of strategies in the $-\infty$ case), by switching between two positional strategies σ_{Min}^* and $\sigma_{\text{Min}}^\dagger$: σ_{Min}^* accumulates negative weight by following negative cycles, and $\sigma_{\text{Min}}^\dagger$ ensures reaching a target. The optimal strategy of Min follows the decisions of σ_{Min}^* , until switching to the decisions of $\sigma_{\text{Min}}^\dagger$ when the length of the play is greater than a finite bound k . These strategies thus require finite memory in the form of a counter.

An interesting case happens if \mathcal{G} has no cycles of negative cumulative weight, *e.g.* if weights are non-negative.

Lemma 11.2. *If \mathcal{G} has no cycles of negative cumulative weight, then both players have optimal strategies that are positional. Moreover, $\text{Val}_{\mathcal{G}} = \text{Val}_{\mathcal{G}}^{|S|}$ and the optimal strategies can be computed in polynomial time.*

Proof. Any strategy of Min is optimal on states of value $+\infty$, so we will ignore those. As there are no negative cycles, value $-\infty$ can only be obtained through reaching a target with final weight $-\infty$. As a consequence, Min has an optimal strategy σ_{Min} that switches between σ_{Min}^* and $\sigma_{\text{Min}}^\dagger$ when the length of the play is greater than some k , as detailed in [BGHM16]. The strategy σ_{Min}^* is only compatible with cycles of negative cumulative weight, and $k \geq |S| + 1$, therefore $\sigma_{\text{Min}}^\dagger$ is never used. Thus, Min has an optimal strategy σ_{Min}^* that is positional. Then, consider $\text{Val}_{\mathcal{G}}^{|S|}$, the value with bounded horizon $|S|$. For every state s , we have $\text{Val}_{\mathcal{G}}^{|S|}(s, \sigma_{\text{Min}}^*) = \text{Val}_{\mathcal{G}}(s)$, so that $\text{Val}_{\mathcal{G}}^{|S|} \leq \text{Val}_{\mathcal{G}}$. As $\text{Val}_{\mathcal{G}} \leq \text{Val}_{\mathcal{G}}^k$ holds for any $k \geq 0$, $\text{Val}_{\mathcal{G}} = \text{Val}_{\mathcal{G}}^{|S|}$. Finally, since value iteration converges in $|S|$ steps, the computation of optimal strategies described in [BGHM16] runs in polynomial time. \square

11.2. **Divergent and almost-divergent weighted games.** Our contribution on finite weighted games is to solve in polynomial time the value problem, for a subclass of finite weighted games. This class corresponds to the games that are *almost-divergent* when seen as timed games with zero clocks. To the best of our knowledge, this is the first attempt to solve a non-trivial class of weighted games with arbitrary weights in polynomial time.

Let us first define the class of divergent weighted games in the untimed setting:

Definition 11.3. A weighted game \mathcal{G} is divergent if every cycle ρ of \mathcal{G} satisfies $\text{wt}_{\Sigma}(\rho) \neq 0$.

In particular, if \mathcal{G} is seen as a weighted timed game with no clocks, then $\mathcal{R}(\mathcal{G})$ is isomorphic to \mathcal{G} , so that divergent weighted games are divergent weighted timed games. We will obtain the following results:

Theorem 11.4. *The value problem over finite divergent weighted games is PTIME-complete. Moreover, deciding if a given finite weighted game is divergent is an NL-complete problem when weights are encoded in unary, and is in PTIME when they are encoded in binary.*

With divergent weighted games, we described a class where the value problem is polynomial instead of pseudo-polynomial. This gain in complexity came at a cost, the absence of cycles of weight 0. We argue that some of those cycles can be allowed, and apply the almost-divergent notion to the untimed setting.

Definition 11.5. A weighted game \mathcal{G} is *almost-divergent* if every 0-cycle ρ of \mathcal{G} satisfies the following property: for every decomposition of ρ into smaller cycles ρ' and ρ'' , ρ' and ρ'' are 0-cycles.

We will obtain the following results, extending Theorem 11.4:

Theorem 11.6. *The value problem over finite almost-divergent weighted games is PTIME-complete. Moreover, deciding if a given finite weighted game is almost-divergent is an NL-complete problem when weights are encoded in unary, and is in PTIME when they are encoded in binary.*

We prove Theorems 11.4 and 11.6 at the same time, by proving a PTIME upper bound for the value problem on almost-divergent game, a PTIME lower bound on divergent games, and finally by studying the decision problems associated with membership to the divergent and almost-divergent classes.

11.2.1. Polynomial upper bound. Our algorithm solving almost-divergent games in polynomial time follows the semi-unfolding scheme, that we detailed in the timed setting as a proof of Theorem 3.9. The SCC characterisation of almost-divergence and the notion of kernel apply, and the only changes concern the complexity analysis. In particular, with 0 clocks, we have $|\mathcal{R}(\mathcal{G})| = |\mathcal{G}|$, therefore computing the kernel and the states of value $+\infty$ or $-\infty$ can be done in polynomial time. Moreover, the semi-unfolding of Section 6 has polynomial depth.

In order to compute the value of a state s_0 of \mathcal{G} , one could construct the semi-unfolding of root s_0 , and compute its value. Indeed, every cycle in $\mathcal{T}(\mathcal{G})$ belongs to \mathbf{K} , so they must all be 0-cycles, therefore by Lemma 11.2 we can compute $\text{Val}_{\mathcal{T}(\mathcal{G})}(\tilde{s}_0)$ in time polynomial in the size of $\mathcal{T}(\mathcal{G})$. However, this would be an exponential time algorithm, since the number of nodes in T can be exponential in $|S|$. We argue that this exponential blow-up can be avoided: when two nodes of T are at the same depth and are labelled by the same state they can be merged, producing a graph T that is acyclic instead of tree-shaped, with at most quadratically many states. This does not change the value of the resulting weighted game $\mathcal{T}(\mathcal{G})$ at its root, because the two merged nodes had the same sub-tree, and therefore were states with the same value in $\mathcal{T}(\mathcal{G})$. This optimization on the construction of $\mathcal{T}(\mathcal{G})$ is performed on-the-fly, while the semi-unfolding is constructed, such that constructing $\mathcal{T}(\mathcal{G})$ (and solving it by Lemma 11.2) can be done in time polynomial in the size of \mathcal{G} .

11.2.2. *Polynomial lower bound.* Let us show that the value problem is PTIME-hard on divergent weighted games. This comes from a reduction (in logarithmic space) of the problem of solving finite games with reachability objectives [Imm81]. To a reachability game, we simply set the weight of every transition to 1 and the final weight of every target to 0, making it a divergent weighted game. Then, Min wins the reachability game if and only if the value in the weighted game is lower than $|S|$. The same reduction can be used to show the PTIME-hardness of the $+\infty$ and $-\infty$ -value problems on divergent weighted games.

11.2.3. *Deciding divergence and almost-divergence.* Let us study the *membership problem* for divergent and almost-divergent weighted games, *i.e.* the decision problem that asks if a given weighted game is divergent or almost-divergent, and prove the results of Theorems 11.4 and 11.6.

We start with almost-divergent games, and will rely on the characterization of almost-divergent games in term of SCCs given in Proposition 4.4. First, we note that simple cycles are enough to ensure that an SCC is non-negative (respectively non-positive), providing us with an efficient way to check this property:

Lemma 11.7. *An SCC S is non-negative (respectively non-positive) if and only if every simple cycle in S is non-negative (respectively non-positive). Moreover, deciding if an SCC is non-negative (respectively non-positive) is in NL when weights are encoded in unary, and is in PTIME when they are encoded in binary.*

Proof. The direct implication holds by definition. Reciprocally, let us assume that every simple cycle in S is non-negative (respectively non-positive), and prove that every cycle ρ in S is non-negative (respectively non-positive). The cycle ρ can be decomposed into simple cycles, all belonging to S . Therefore they are all non-negative (respectively non-positive). As the cumulative weight of ρ is the sum of the cumulative weights of these simple cycles, ρ must be non-negative (respectively non-positive).

As a corollary, an SCC S is non-negative (respectively non-positive) if and only if every cycle in S , of length at most $|S|$, is non-negative (respectively non-positive).

To decide if a strongly connected \mathcal{G} is non-negative (respectively non-positive), we outline two procedures: one is deterministic and will provide the polynomial upper bound on time-complexity, the other will guess a logarithmic number of bits and provide NL membership.

The deterministic algorithm proceeds as follows: with Floyd-Warshall's algorithm, one can compute the shortest paths (respectively greatest paths) adjacency matrix M in cubic time, such that $M(s, s')$ contains the minimal (respectively maximal) value in the finite set $\{\text{wt}(\rho) \mid \rho \text{ simple path from } s \text{ to } s'\}$. Then, S is non-negative (respectively non-positive) if and only if for every state s it holds that $M(s, s) \geq 0$ (respectively $M(s, s) \leq 0$). If there exists a state s such that $M(s, s) < 0$ (respectively $M(s, s) > 0$), then S is not non-negative (respectively non-positive) as there is a negative (respectively positive) cycle. Conversely, if $M(s, s) \geq 0$ (respectively $M(s, s) \leq 0$), we know that all simple paths from s to s have non-negative (respectively non-positive) weight.

Let us now assume that weights are encoded in unary, and present a non-deterministic procedure. Then, note that a (binary) register containing integer values in $[-B, B]$, with B polynomial in w_{\max} and $|S|$, requires a number of bits at most logarithmic in the size of \mathcal{G} . An SCC is *not* non-negative (respectively not non-positive), if and only if it contains a cycle of positive (respectively negative) cumulative weight, of length bounded by $|S|$. We

can guess such a cycle ρ on-the-fly, keeping in memory its cumulative weight (smaller than $B = w_{\max} \times |S|$ in absolute value), its initial state, and its current length, all in logarithmic space. If the length of the cycle exceeds $|S|$, the guess is invalid. Similarly, we can verify that the last state equals the first, and that the computed cumulative weight is indeed positive (respectively negative). Therefore, deciding if S is non-negative (respectively not non-positive) is in $\text{coNL} = \text{NL}$ [Imm88, Sze88]. Note that when weights are encoded in binary this procedure only gives coNP membership. \square

Let us now explain why the membership problem is an NL -complete problem when weights are encoded in unary. First, to prove the membership in NL , notice that a weighted game is *not almost-divergent* if and only if there is a positive cycle and a negative cycle, both of length at most $|S|$, and belonging to the same SCC. This can be tested in NL , using a non-deterministic procedure similar to the one from Lemma 11.7. We first guess a starting state for both cycles. Verifying that those are in the same SCC can be done in NL by using standard reachability analysis. Then, we once again guess the two cycles on-the-fly, keeping in memory their cumulative weights in logarithmic space. Therefore, testing divergence is in $\text{coNL} = \text{NL}$ [Imm88, Sze88].

The NL -hardness (indeed coNL -hardness, which is equivalent [Imm88, Sze88]) is shown by a reduction from the reachability problem in a finite graph. More precisely, we consider a finite automaton with a starting state and a different target state without outgoing transitions. We construct from it a weighted game by distributing all states to Min , and equipping all transitions with weight 0. We also add a loop with weight 1 on the initial state, one with weight -1 on the target state, and a transition from the target state to the initial state with weight 0. Then, the game is not almost-divergent if and only if the target can be reached from the initial state in the automaton.

When weights are encoded in binary, the previous decision procedure gives NP membership. However, we can achieve a PTIME upperbound by computing the strongly connected components and then using Lemma 11.7 to check that each SCC is either non-negative or non-positive.

This concludes the proofs of Theorem 11.6. For the membership results in Theorem 11.4, we note that a weighted game \mathcal{G} is divergent if and only if it is almost-divergent and its kernel is empty, *i.e.* \mathcal{G} does not contain any simple cycle of weight zero. Variants of the previous techniques can be used to either compute the kernel in polynomial time, guess a 0-cycle of length at most $|S|$ in NL when weights are encoded in unary, or prove NL -hardness.

12. CONCLUSION

We have obtained new results for several controller synthesis problems on timed automata, that we now summarise. Our study of weighted timed games belongs to a series of works that explore the frontier of decidability. We introduced the first decidable class of weighted timed games with arbitrary weights and no restrictions on the number of clocks. We have given an approximation procedure for a larger class of weighted timed games, where the exact problem becomes undecidable. In addition, we have proved the correctness of a symbolic approximation schema, that does not start by splitting exponentially every region, but only does so when necessary. We argue that this paves the way towards an implementation of value approximation for weighted timed games. Such tool would likely struggle with instances

of moderate size, but could help with the design and testing of alternative approaches that trade theoretical guarantees with performance.

Another perspective is to extend this work to the concurrent setting, where both players play simultaneously and the shortest delay is selected. It should be noted that several known results on weighted timed games with non-negative weights [BCFL04, ABM04, BJM15] are stated in such a concurrent setting. We did not consider this setting in this work because concurrent WTGs are not determined, and several of our proofs rely on this property for symmetrical arguments (mainly to lift results of non-negative strongly connected components to non-positive ones).

A long-standing open problem is the approximation of weighted timed games, *i.e.* whether one can compute an arbitrarily close approximation of the value of a given game. We successfully solved this problem on the class of almost-divergent games, but we were not able to extend further our techniques to more general games. As a first step, we could try to consider the slightly larger class of 0-isolated games, where we ask for every cycle of the region game to have a weight either ≥ 1 , or ≤ -1 , or exactly 0. We do not have approximation results on this 0-isolated class, and as such it forms a natural intermediate step between the best known decidable class and the general case. However we must prepare ourselves to possibly negative answers: the value of a weighted timed game could be *non approximable*, though we are not aware of any such game; especially we do not even know what kind of real numbers (irrational, or even transcendental) could be obtained as the value in the null valuation. Therefore, pursuing better lower bounds in various settings could help in the future, in order to close the remaining complexity gaps.

The divergence and almost-divergence classes that we have studied in this article are independent of the partition of locations into players. It would be interesting, as future work, to investigate restrictions taking into account the interaction of the two players.

REFERENCES

- [ABM04] 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.
- [AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
- [ALTP04] Rajeev Alur, Salvatore La Torre, and George J. Pappas. Optimal paths in weighted timed automata. *Theoretical Computer Science*, 318(3):297–322, 2004.
- [AM99] Eugene Asarin and Oded Maler. As soon as possible: Time optimal control for timed automata. In *Hybrid Systems: Computation and Control*, volume 1569 of *Lecture Notes in Computer Science*, pages 19–30. Springer, 1999.
- [AMPS98] Eugene Asarin, Oded Maler, Amir Pnueli, and Joseph Sifakis. Controller synthesis for timed automata. *IFAC Proceedings Volumes*, 31(18):447–452, 1998.
- [BBBR07] 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.
- [BBL08] Patricia Bouyer, Ed Brinksma, and Kim G. Larsen. Optimal infinite scheduling for multi-priced timed automata. *Formal Methods in System Design*, 32(1):3–23, 2008.
- [BBM06] Patricia Bouyer, Thomas Brihaye, and Nicolas Markey. Improved undecidability results on weighted timed automata. *Information Processing Letters*, 98(5):188–194, 2006.

- [BBR05] Thomas Brihaye, Véronique Bruyère, and Jean-François Raskin. On optimal timed strategies. In *Proceedings of the Third international conference on Formal Modeling and Analysis of Timed Systems (FORMATS'05)*, volume 3829 of *LNCS*, pages 49–64. Springer, 2005.
- [BCFL04] Patricia Bouyer, Franck Cassez, Emmanuel Fleury, and Kim G. Larsen. Optimal strategies in priced timed game automata. In *Proceedings of the 24th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'04)*, volume 3328 of *LNCS*, pages 148–160. Springer, 2004.
- [BCR14] Romain Brenguier, Franck Cassez, and Jean-François Raskin. Energy and mean-payoff timed games. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control (HSCC'14)*, pages 283–292. ACM, 2014.
- [BDM⁺98] Marius Bozga, Conrado Daws, Oded Maler, Alfredo Olivero, Stavros Tripakis, and Sergio Yovine. Kronos: A model-checking tool for real-time systems. In Alan J. Hu and Moshe Y. Vardi, editors, *Computer Aided Verification (CAV 1998), Proceedings*, pages 546–550, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.
- [BFH⁺01] Gerd Behrmann, Ansgar Fehnker, Thomas Hune, Kim G. Larsen, Judi Romijn, and Frits W. Vaandrager. Minimum-cost reachability for priced timed automata. In *Proceedings of the 4th International Workshop on Hybrid Systems: Computation and Control (HSCC'01)*, volume 2034 of *Lecture Notes in Computer Science*, pages 147–161. Springer, 2001.
- [BG19] Damien Busatto-Gaston. *Symbolic controller synthesis for timed systems: robustness and optimality*. Theses, Aix Marseille Université, December 2019. URL: <https://hal.archives-ouvertes.fr/tel-02436831>.
- [BGH⁺15] 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.
- [BGHM16] Thomas Brihaye, Gilles Geeraerts, Axel Haddad, and Benjamin Monmege. Pseudopolynomial iterative algorithm to solve total-payoff games and min-cost reachability games. *Acta Informatica*, 2016. doi:10.1007/s00236-016-0276-z.
- [BGMR17] Damien Busatto-Gaston, Benjamin Monmege, and Pierre-Alain Reynier. Optimal reachability in divergent weighted timed games. In Javier Esparza and Andrzej S. Murawski, editors, *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, Uppsala, Sweden, April 2017. Springer. doi:10.1007/978-3-662-54458-7_10.
- [BGMR18] Damien Busatto-Gaston, Benjamin Monmege, and Pierre-Alain Reynier. Symbolic Approximation of Weighted Timed Games. In Sumit Ganguly and Paritosh Pandya, editors, *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2018)*, volume 122 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 28:1–28:16, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2018/9927>, doi:10.4230/LIPIcs.FSTTCS.2018.28.
- [BGNK⁺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, pages 560–575. Springer, 2014. doi:10.1007/978-3-662-44584-6_38.
- [BJM15] 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 *Leibniz International Proceedings in Informatics*, pages 311–324. Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.CONCUR.2015.311.
- [BL69] J. Richard Büchi and Lawrence H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.
- [BPDG98] Béatrice Bérard, Antoine Petit, Volker Diekert, and Paul Gastin. Characterization of the expressive power of silent transitions in timed automata. *Fundamenta Informaticae*, 36(2-3):145–182, 1998.
- [HIJM13] 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)*, volume 8052 of *Lecture Notes in Computer Science*, pages 531–545. Springer, 2013.
- [HPT19] Frédéric Herbreteau, Gerald Point, and Thanh-Tung Tran. Tchecker. <http://www.labri.fr/perso/herbrete/tchecker/index.html>, 2019.
- [HSW10] Frédéric Herbreteau, B. Srivathsan, and Igor Walukiewicz. Efficient emptiness check for timed Büchi automata. In *Computer Aided Verification (CAV 2010), Proceedings*, volume 6174 of *Lecture Notes in Computer Science*, pages 148–161. Springer, 2010. doi:10.1007/978-3-642-14295-6_15.
- [Imm81] Neil Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22(3):384–406, 1981.
- [Imm88] Neil Immerman. Nondeterministic space is closed under complementation. *SIAM Journal on Computing*, 17:935–938, 1988.
- [JT07] Marcin Jurdziński 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 *LNCS*, pages 838–849. Springer, 2007.
- [KBB⁺08] Leonid Khachiyan, Endre Boros, Konrad Borys, Khaled Elbassioni, Vladimir Gurvich, Gabor Rudolf, and Jihui Zhao. On short paths interdiction problems: Total and node-wise limited interdiction. *Theory of Computing Systems*, 43(2):204–233, 2008. doi:10.1007/s00224-007-9025-6.
- [LBB⁺01] Kim G. Larsen, Gerd Behrmann, Ed Brinksma, Ansgar Fehnker, Thomas Hune, Paul Pettersson, and Judi Romijn. As cheap as possible: Efficient cost-optimal reachability for priced timed automata. In *Proceedings of the 13th International Conference on Computer Aided Verification (CAV'01)*, volume 2102 of *Lecture Notes in Computer Science*, pages 493–505. Springer, 2001.
- [LPY97] Kim G. Larsen, Paul Pettersson, and Wang Yi. Uppaal in a nutshell. *International Journal on Software Tools for Technology Transfer*, 1(1):134–152, Dec 1997. doi:10.1007/s100090050010.
- [Mat02] Jiri Matousek. *Lectures on Discrete Geometry*. Springer-Verlag, Berlin, Heidelberg, 2002.
- [Pur00] Anuj Puri. Dynamical properties of timed automata. *Discrete Event Dynamic Systems*, 10(1-2):87–113, 2000.
- [Rut11] Michał Rutkowski. Two-player reachability-price games on single-clock timed automata. In *Proceedings of the Ninth Workshop on Quantitative Aspects of Programming Languages (QAPL'11)*, volume 57 of *Electronic Proceedings in Theoretical Computer Science*, pages 31–46, 2011.
- [Sav70] Walter J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4(2):177–192, 1970.
- [Sze88] Róbert Szelepcsényi. The method of forced enumeration for nondeterministic automata. *Acta Informatica*, 26(3):279–284, 1988.
- [WTH92] Howard Wong-Toi and Gerard Hoffmann. The control of dense real-time discrete event systems. Technical report, DTIC Document, 1992.

APPENDIX A. PROOF OF LEMMA 2.6

The number $\text{Splits}(m, n)$ is known in discrete geometry as the number of faces obtained when partitioning \mathbb{R}^n by m hyperplanes. We refer to [Mat02, Chap. 6.1] for more details on this problem, and make use of their vocabulary in this proof, so that cells are called faces and affine equalities are called hyperplanes. The faces are relatively open convex sets, with dimensions ranging from 0 to n . We call k -faces the faces of dimension k . For example, in Figure 5 the 0-face is the point at the intersection of the two hyperplanes, the four 1-faces are half-lines, and the four 2-faces partition the remaining points of \mathbb{R}^2 .

According to [Mat02, Prop. 6.1.1], the number of n -faces in a space of dimension n is $\sum_{\ell=0}^{\min(n,m)} \binom{m}{\ell}$. These faces are known as faces of full dimension. If n or m equals 0, $\text{Splits}(m, n) = 1$. If $n, m \geq 1$, the $(n-1)$ -faces are obtained in each of our hyperplanes (spaces of dimension $n-1$), by considering their partition by the $m-1$ other hyperplanes. In these partitions, the $(n-1)$ -faces are of full dimension, so that each hyperplane contains $\sum_{\ell=0}^{\min(n-1,m-1)} \binom{m-1}{\ell}$ such faces. The total number of $(n-1)$ -faces is therefore $m \sum_{\ell=0}^{\min(n-1,m-1)} \binom{m-1}{\ell}$. Let us now count the number of k -faces for $k \in [0, n-2]$.

As explained in [Mat02, Chap. 6.1], the number of faces is maximised when the hyperplanes are in general position, *i.e.* for any $k \in [2, \min(n+1, m)]$, the intersection of any set of k hyperplanes is $(n-k)$ -dimensional.²² Then, if $n \geq 2$ the faces of dimension $k \in [\max(0, n-m), n-2]$ lie at the intersection of $n-k$ hyperplanes, and for each such intersection L the k -faces are obtained by partitioning L by the $m-n+k$ other hyperplanes. They are faces of full dimension in L , so that L contains $\sum_{\ell=0}^{\min(k, m-n+k)} \binom{m-n+k}{\ell}$ such faces. There are $\binom{m}{n-k}$ intersections L , so that the total number of k -faces is $\sum_{\ell=0}^{\min(k, m-n+k)} \binom{m}{n-k} \binom{m-n+k}{\ell}$. Note that this formula matches the number of $(n-1)$ -faces and the number of n -faces previously given, and that if $m \leq n$ there are no faces of dimension $k < n-m$. Therefore,

$$\text{Splits}(m, n) = \sum_{k=\max(0, n-m)}^n \sum_{\ell=0}^{\min(k, m-n+k)} \binom{m}{n-k} \binom{m-n+k}{\ell}.$$

Let us show $\text{Splits}(m, n) \leq 2^n(m+1)^n$. First, note that for all $0 \leq \ell \leq k \leq n$ with $n-k \leq m$ and $\ell \leq m-n+k$, it holds that $\binom{m}{n-k} \binom{m-n+k}{\ell} = \binom{m}{n-k+\ell} \binom{n-k+\ell}{n-k} \leq \binom{m}{n-k+\ell} \binom{n}{k}$. Then, $\text{Splits}(m, n) \leq \sum_{k=\max(0, n-m)}^n \left[\binom{n}{k} \sum_{\ell=0}^{\min(k, m-n+k)} \binom{m}{n-k+\ell} \right]$. Moreover, $\sum_{\ell=0}^{\min(k, m-n+k)} \binom{m}{n-k+\ell} \leq \sum_{\ell=0}^{\min(m, n)} \binom{m}{\ell}$.

Then, note that $\sum_{\ell=0}^n \binom{m}{\ell}$ is the number of selections of size at most n out of a set of m elements, without repetitions. It is smaller than the number of selections of n elements, with repetitions, in a set of size $m+1$ (the extra element can be used as padding). Therefore, $\sum_{\ell=0}^{\min(n, m)} \binom{m}{\ell} \leq (m+1)^n$. On the other hand, $\sum_{k=\max(0, n-m)}^n \binom{n}{k} \leq \sum_{k=0}^n \binom{n}{k} = 2^n$. It follows that $\text{Splits}(m, n) \leq 2^n(m+1)^n$.

APPENDIX B. EXAMPLE OF AN EXECUTION OF THE APPROXIMATION SCHEMA

Consider the WTG \mathcal{G} in Figure 10 and some $\varepsilon > 0$. We want to compute an ε -approximation of its value in location ℓ_0 for the valuation $(x_1=0, x_2=0)$, denoted $\text{Val}_{\mathcal{G}}(\ell_0, (0, 0))$. In this example, we will use $\varepsilon = 15$ because the computations would not be readable with a smaller

²²for $n-k < 0$ this means that the intersection is empty.

precision. Since in this example each location ℓ of \mathcal{G} leads to a unique state (ℓ, r) of $\mathcal{R}(\mathcal{G})$, we will refer to states of $\mathcal{R}(\mathcal{G})$ by their associated location label. As explained in Example 3.8, $\mathcal{R}(\mathcal{G})$ contains one SCC made of two simple cycles, $\mathbf{p}_1 = \ell_1 \rightarrow \ell_2 \rightarrow \ell_1$ is a positive cycle and $\mathbf{p}_2 = \ell_1 \rightarrow \ell_3 \rightarrow \ell_4 \rightarrow \ell_1$ is a 0-cycle.

Therefore, $\mathcal{R}(\mathcal{G})$ only contains non-negative SCCs and is almost-divergent. Since all states are in the attractor of Min towards L_t , all cycles are non-negative and the final weight function is bounded (on all reachable regions), there are no configurations in $\mathcal{R}(\mathcal{G})$ with value $+\infty$ or $-\infty$.

We let the kernel K be the sub-game of $\mathcal{R}(\mathcal{G})$ defined by \mathbf{p}_2 , and we construct a semi-unfolding $\mathcal{T}(\mathcal{G})$ of $\mathcal{R}(\mathcal{G})$ of equivalent value. We should unfold the game until every stopped branch contains a state seen at least $3|\mathcal{R}(\mathcal{G})|w_{\max} + 2w_{\max}^t + 2 = 3 \times 3 \times 4 + 2 \times 1 + 2 = 40$ times. We will unfold with bound 4 instead of 40 for readability (it is enough on this example). Thus the infinite branch $(\ell_1 \ell_2)^\omega$ is stopped when ℓ_1 is reached for the fourth time, as depicted in Figure 18.

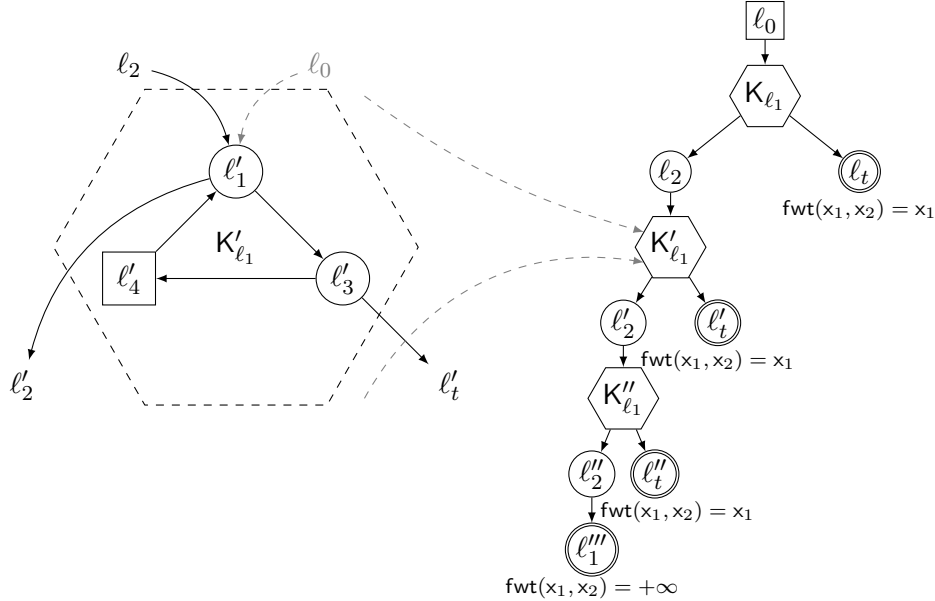


Figure 18: The kernel K (with input state ℓ_1), and a semi-unfolding $\mathcal{T}(\mathcal{G})$ such that $\text{Val}_{\mathcal{G}}(\ell_0, (0, 0)) = \text{Val}_{\mathcal{T}(\mathcal{G})}(\ell_0, (0, 0))$. We denote by ℓ_i , ℓ'_i and ℓ''_i the locations in K , K' and K'' .

Let us now compute an approximation of $\text{Val}_{\mathcal{T}(\mathcal{G})}$. Let us first remove the states of value $+\infty$: ℓ'''_1 and ℓ''_2 . Then, we start at the bottom and compute an $(\varepsilon/3)$ -approximation of the value of ℓ''_1 in the game defined by K''_{ℓ_1} and its output edge to ℓ''_1 . Following Section 8.2, we should use $N \geq 3(4 + 1)/\varepsilon$ and compute values in the $1/N$ -corners game $\mathcal{C}_N(K''_{\ell_1})$ in order to obtain an $(\varepsilon/3)$ -approximation of the value function. For $\varepsilon = 15$ we will use $N = 1$ (in this case the computation happens to be exact and would also hold with a small ε). We construct this corner game, and obtain the finite (untimed) weighted game in Figure 19.

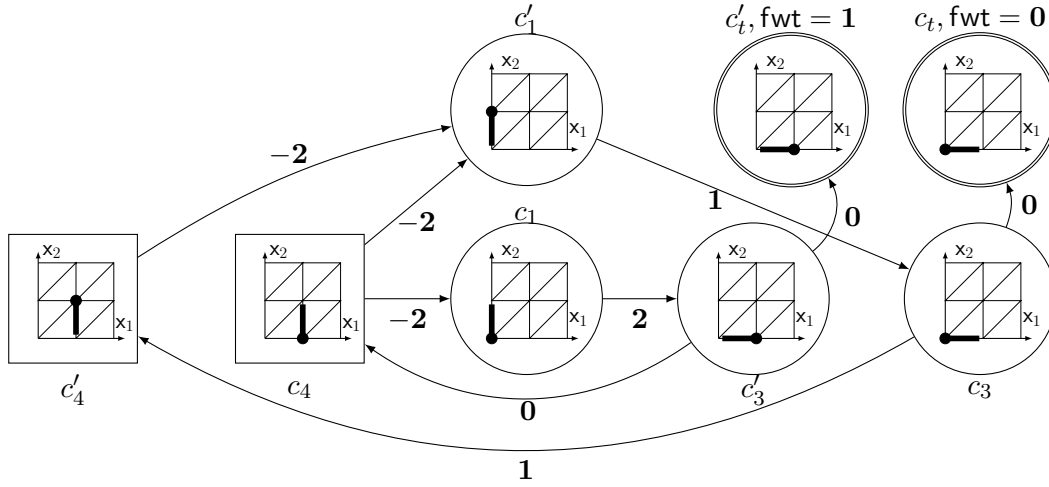


Figure 19: The finite weighted game obtained from $\mathcal{C}_1(K''_{\ell'_1})$, where c_i and c'_i are the corners of ℓ''_i in $\mathcal{T}(\mathcal{G})$.

We can compute the values in this game to obtain $\text{Val}(c'_1) = 1$ and $\text{Val}(c_1) = 3$. We then define a value for every configuration in state ℓ''_1 by linear interpolation, obtaining:

$$(x_1, x_2) \mapsto 3 - 2x_2 .$$

This happens to be exactly $(x_1, x_2) \mapsto \text{Val}_{\mathcal{T}(\mathcal{G})}(\ell''_1, (x_1, x_2))$ in this case, but would only be an $\varepsilon/3$ -approximation of it in general. Now, we can compute an $\varepsilon/3$ -approximation of $\text{Val}_{\mathcal{T}(\mathcal{G})}(\ell'_1)$ with one step of value iteration, obtaining

$$(x_1, x_2) \mapsto \inf_{0 < d < 2 - x_1} (-1) \times d + 1 + 3 - 2(0 + d) = 3x_1 - 2 .$$

The next step is computing an $\varepsilon/3$ -approximation of the value of ℓ'_1 in the game defined by $K'_{\ell'_1}$ and its output edges to ℓ'_t and ℓ'_2 , of respective final weight functions $(x_1, x_2) \mapsto x_1$ and $(x_1, x_2) \mapsto 3x_1 - 2$. This will give us a $2\varepsilon/3$ -approximation of $\text{Val}_{\mathcal{T}(\mathcal{G})}(\ell'_1)$.

Following Section 8.2 once again, we should use $N \geq 3(5 + 3)/\varepsilon$ and compute values in the $1/N$ -corners game $\mathcal{C}_N(K'_{\ell'_1})$. For $\varepsilon = 15$ this gives $N = 2$ (which will once again keep the computation exact). We can construct a finite (untimed) weighted game as in Figure 19, and obtain a value for each $1/2$ -corner of state ℓ'_1 :

- On the $1/2$ -region $(x_1 = 0, 0 < x_2 < 1/2)$, corner $(0, 0)$ has value 2 and corner $(0, 1/2)$ has value 2.
- On the $1/2$ -region $(x_1 = 0, x_2 = 1/2)$, corner $(0, 1/2)$ has value 2.
- On the $1/2$ -region $(x_1 = 0, 1/2 < x_2 < 1)$, corner $(0, 1/2)$ has value 2 and corner $(0, 1)$ has value 1.

From these results, we define a piecewise affine function by interpolating the values of corners on each $1/2$ -region, and obtain

$$(x_1, x_2) \mapsto \begin{cases} 2 & \text{if } x_2 \leq 1/2 \\ 3 - 2x_2 & \text{otherwise} \end{cases}$$

as depicted in Figure 20.

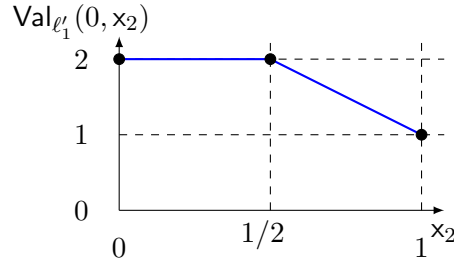


Figure 20: The value function $(x_1, x_2) \mapsto \text{Val}_{\mathcal{T}(\mathcal{G})}(\ell'_1, (x_1, x_2))$, projected on $x_1 = 0$. Black dots represent the values obtained for $1/2$ -corners using the corner-point abstraction.

This gives us a $2\varepsilon/3$ -approximation of $(x_1, x_2) \mapsto \text{Val}_{\mathcal{T}(\mathcal{G})}(\ell'_1, (x_1, x_2))$ (in fact exactly $\text{Val}_{\mathcal{T}(\mathcal{G})}(\ell'_1)$). Now, we can compute a $2\varepsilon/3$ -approximation of $\text{Val}_{\mathcal{T}(\mathcal{G})}(\ell_2)$ on region $(1 < x_1 < 2, x_2 = 0)$ with one step of value iteration, obtaining :

$$(x_1, x_2) \rightarrow \inf_{0 < d < 2 - x_1} \begin{cases} 3 - d & \text{if } d \leq 1/2 \\ 4 - 3d & \text{otherwise} \end{cases} = \begin{cases} 3x_1 - 2 & \text{if } x_1 \leq 3/2 \\ x_1 + 1 & \text{otherwise} \end{cases}$$

Then, we need to compute an $\varepsilon/3$ -approximation of the value of ℓ_1 in the game defined by \mathcal{K}_{ℓ_1} and its output edges to ℓ_t and ℓ_2 , of respective final weight functions $(x_1, x_2) \mapsto x_1$ and $(x_1, x_2) \mapsto 3x_1 - 2$ if $x_1 \leq 3/2$; $x_1 + 1$ otherwise. This will give us an ε -approximation of $\text{Val}_{\mathcal{T}(\mathcal{G})}(\ell_1)$.

Following Section 8.2 one last time, we should use $N \geq 3(5 + 3)/\varepsilon$ and compute values in the $1/N$ -corner game $\mathcal{C}_N(\mathcal{K}_{\ell_1})$. This time, let us use $N = 3$ to showcase an example where the computed value is not exact. We can construct a finite (untimed) weighted game as in Figure 19, and obtain a value for each $1/3$ -corner of state ℓ'_1 . From these results, we define a piecewise affine function by interpolation, as depicted in Figure 21.

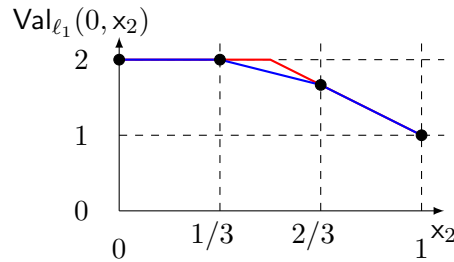


Figure 21: The value function $(x_1, x_2) \mapsto \text{Val}_{\mathcal{T}(\mathcal{G})}(\ell_1, (x_1, x_2))$, projected on $x_1 = 0$, is depicted in red. Black dots represent the values obtained for $1/3$ -corners using the corner-points abstraction, and the derived approximation of the value function is depicted in blue

Finally, from this ε -approximation of $\text{Val}_{\mathcal{T}(\mathcal{G})}(\ell_1)$, we can compute an ε -approximation of $\text{Val}_{\mathcal{T}(\mathcal{G})}(\ell_0)$ using one step of value iteration, and conclude. On our example this ensures

$$\text{Val}_{\mathcal{T}(\mathcal{G})}(\ell_0, (0, 0)) = \sup_{0 < d < 1} \text{Val}_{\mathcal{T}(\mathcal{G})}(\ell_1, (0, d)) \in [2 - \varepsilon, 2 + \varepsilon].$$