



HAL
open science

One-Clock Priced Timed Games with Arbitrary Weights

Thomas Brihaye, Gilles Geeraerts, Axel Haddad, Engel Lefauchaux, Benjamin Monmege

► **To cite this version:**

Thomas Brihaye, Gilles Geeraerts, Axel Haddad, Engel Lefauchaux, Benjamin Monmege. One-Clock Priced Timed Games with Arbitrary Weights. Logical Methods in Computer Science, 2022, 18 (3), pp.51. hal-02424743v2

HAL Id: hal-02424743


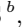


<https://hal.science/hal-02424743v2>

Submitted on 27 Sep 2022

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.

ONE-CLOCK PRICED TIMED GAMES WITH NEGATIVE WEIGHTS

THOMAS BRIHAYE ^a, GILLES GEERAERTS ^b, AXEL HADDAD ^a, ENGEL LEFAUCHEUX ^c,
AND BENJAMIN MONMEGE ^d

^a Université de Mons, Belgium

e-mail address: thomas.brihaye@umons.ac.be, axel.haddad@umons.ac.be

^b Université libre de Bruxelles, Belgium

e-mail address: gigeerae@ulb.ac.be

^c Max-Planck Institute for Software Systems, Germany

e-mail address: elefauch@mpi-sws.org

^d Aix Marseille Univ, LIS, CNRS, Marseille, France

e-mail address: benjamin.monmege@univ-amu.fr

ABSTRACT. Priced timed games are two-player zero-sum games played on priced timed automata (whose locations and transitions are labeled by weights modelling the cost of spending time in a state and executing an action, respectively). The goals of the players are to minimise and maximise the cost to reach a target location, respectively. We consider priced timed games with one clock and arbitrary integer weights and show that, for an important subclass of them (the so-called *simple* priced timed games), one can compute, in pseudo-polynomial time, the optimal values that the players can achieve, with their associated optimal strategies. As side results, we also show that one-clock priced timed games are determined and that we can use our result on simple priced timed games to solve the more general class of so-called *negative-reset-acyclic* priced timed games (with arbitrary integer weights and one clock). The decidability status of the full class of priced timed games with one-clock and arbitrary integer weights still remains open.

1. INTRODUCTION

Game theory is nowadays a well-established framework in theoretical computer science, enabling computer-aided design of computer systems that are correct-by-construction. It allows one to describe and analyse the possible interactions of antagonistic agents (or players) as in the *controller synthesis* problem, for instance. This problem asks, given a model of the environment of a system, and of the possible actions of a controller, to compute a controller

Key words and phrases: Priced timed games; Real-time systems; Game theory.

A preliminary version of this work has been published in the proceedings of FSTTCS 2015 [BGH⁺15]. The research leading to these results was funded by the European Union Seventh Framework Programme (FP7/2007-2013) under Grant Agreement n°601148 (CASSTING). This work was also partly supported by the Fonds de la Recherche Scientifique - FNRS under grant n°T.0027.21. During part of the preparation of this article, the last author was (partially) funded by the ANR project DeLTA (ANR-16-CE40-0007) and the ANR project Ticktac (ANR-18-CE40-0015).




Figure 1: A simple priced timed game (left) and the lower value function of location l_1 (right). Transitions without label have weight 0.

that constraints the environment to respect a given specification. Clearly, one cannot assume in general that the two players (the environment and the controller) will collaborate, hence the need to find a *strategy for the controller* that enforces the specification *whatever the environment does*. This question thus reduces to computing a so-called winning strategy for the corresponding player in the game model.

In order to describe precisely the features of complex computer systems, several game models have been considered in the literature. In this work, we focus on the model of Priced Timed Games (PTGs for short), which can be regarded as an extension (in several directions) of classical finite automata. First, like timed automata [AD94], PTGs have *clocks*, which are real-valued variables whose values evolve with time elapsing, and which can be tested and reset along the transitions. Second, the locations are associated with weights representing rates and transitions are labeled by discrete weights, as in priced timed automata [BFH⁺01, ALTP04, BBR07]. These weights allow one to associate a price with each play (or run), which depends on the sequence of transitions traversed by the play, and on the time spent in each visited location. Finally, a PTG is played by two players, called Min and Max, and each location of the game is owned by either of them (we consider a turn-based version of the game). The player who controls the current location decides how long to wait, and which transition to take.

In this setting, the goal of Min is to reach a given set of target locations, while minimising the price of the play to reach such a location. Player Max has an antagonistic objective: it tries to avoid the target locations, and, if not possible, to maximise the accumulated cost up to the first visit of a target location. To reflect these objectives, we define the upper value $\overline{\text{Val}}$ of the game as a mapping of the configurations of the PTG to the least price that Min can guarantee while reaching the target, whatever the choices of Max. Similarly, the lower value $\underline{\text{Val}}$ returns the greatest price that Max can ensure (letting the price be $+\infty$ in case the target locations are not reached).

An example of PTG is given in Figure 1, where the locations of Min and Max are represented by circles and rectangles respectively. The integers next to the locations are their rates, i.e. the cost of spending one time unit in the location. Moreover, there is only one clock x in the game, which is never reset, and all guards on transitions are $x \in [0, 1]$ which force every player to keep the clock value below or equal to 1, but do not hinder the choice of transition (hence this guard is not displayed and transitions are only labelled by their respective discrete weight): this is an example of a *simple priced timed game* (we will define them properly later). It is easy to check that Min can force reaching the target location l_f from all configurations (ℓ, ν) of the game, where ℓ is a location and ν is a real

value of the clock in $[0, 1]$. Let us comment on the optimal strategies for both players. From a configuration (ℓ_4, ν) , with $\nu \in [0, 1]$, Max better waits until the clock takes value 1, before taking the transition to ℓ_f (it is forced to move, by the rules of the game). Hence, Max's optimal value is $3(1 - \nu) - 7 = -3\nu - 4$ from all configurations (ℓ_4, ν) . Symmetrically, it is easy to check that Min better waits as long as possible in ℓ_7 , hence its optimal value is $-16(1 - \nu)$ from all configurations (ℓ_7, ν) . However, optimal value functions are not always *that simple*, see for instance the lower value function of ℓ_1 on the right of Figure 1, which is a piecewise affine function. To understand why value functions can be piecewise affine, consider the sub-game enclosed in the dotted rectangle in Figure 1, and consider the value that Min can guarantee from a configuration of the form (ℓ_3, ν) in this sub-game. Clearly, Min must decide how long it will spend in ℓ_3 and whether it will go to ℓ_4 or ℓ_7 . Its optimal value from all (ℓ_3, ν) is thus

$$\inf_{0 \leq t \leq 1 - \nu} \min(4t + 3(1 - (\nu + t)) - 7, 4t + 6 - 16(1 - (\nu + t))) = \min(-3\nu - 4, 16\nu - 10).$$

Since $16\nu - 10 \geq -3\nu - 4$ if and only if $\nu \geq 6/19$, an optimal choice of Min is to move instantaneously to ℓ_7 if $\nu \in [0, 6/19]$ and to move instantaneously to ℓ_4 if $\nu \in (6/19, 1]$, hence the value function of ℓ_3 (in the sub-game) is a piecewise affine function with two pieces.

Related work. PTGs are a special case of hybrid games [dAHM01, MPS95, WT97], independently investigated in [BCFL04] and [ABM04]. For (non-necessarily turn-based) PTGs with *non-negative* weights, semi-algorithms are given to decide the *value problem*, that is to say, whether the upper value of a location (the best price that Min can guarantee starting with a clock value 0), is below a given threshold. It was also shown that, under the *strongly non-Zeno assumption* on weights (asking the existence of $\kappa > 0$ such that every cycle in the underlying region graph has a weight at least κ), the proposed semi-algorithms always terminate. This assumption was justified in [BBR05, BBM06] by showing that, without it, the *existence problem*, that is to decide whether Min has a strategy guaranteeing to reach a target location with a price below a given threshold, is indeed undecidable for PTGs with non-negative weights and three or more clocks. This result was recently extended in [BJM14] to show that the value problem is also undecidable for PTGs with non-negative weights and four or more clocks. In [BCJ09], the undecidability of the existence problem has also been shown for PTGs with arbitrary weights on locations (without weights on transitions), and two or more clocks. Finally, PSPACE-hardness of the value problem has been established for one-clock PTGs in [FIJS20]. On a positive side, the value problem was shown decidable by [BLMR06] for PTGs with one clock when the weights are non-negative: a 3-exponential time algorithm was first proposed, further refined in [Rut11, HIJM13] into an exponential time algorithm. The key point of those algorithms is to reduce the problem to the computation of optimal values in a restricted family of PTGs called *Simple Priced Timed Games* (SPTGs for short), where the underlying automata contain no guard, no reset, and the play is forced to stop after one time unit. More precisely, the PTG is decomposed into a sequence of SPTGs whose value functions are computed and re-assembled to yield the value function of the original PTG. Alternatively, and with radically different techniques, a pseudo-polynomial time algorithm to solve one-clock PTGs with arbitrary weights on transitions, and rates restricted to two values amongst $\{-d, 0, +d\}$ (with $d \in \mathbb{N}$) was given in [BGK⁺14]. More recently, a large subclass of PTGs with arbitrary weights and no restrictions on the number of clocks was introduced in [BGMR17], whose value can be computed in double-exponential time: they are defined via a partition of strongly connected

components with respect to the sign of all the cycles they contain. A survey summarising results on PTGs can be found in [Bou15].

Contributions. Following the decidability results sketched above, we consider PTGs with one clock. We extend those results by considering arbitrary (positive and negative) weights. Indeed, all previous works on PTGs with only one clock (except [BGK⁺14]) have considered non-negative weights only, and the status of the more general case with arbitrary weights has so far remained elusive. Yet, arbitrary weights are an important modelling feature. Consider, for instance, a system which can consume but also produce energy at different rates. In this case, energy consumption could be modelled as a positive rate, and production by a negative rate. As another example, imagine the billing system for electrical power, in a smart house that itself produces energy: the money spent/earned while using/producing energy contains both timed components (the more energy the house produces, the more money the owner gets), and discrete ones (the electricity provider charges discrete costs per month). Such a model has been studied in [BDG⁺16]. In the untimed setting, such extension to negative weights has been considered in [BGHM15, BGHM16]: our result heavily builds upon techniques investigated in these works, as we will see later. Our main contribution is a *pseudo-polynomial time algorithm to compute the value of one-clock SPTGs with arbitrary weights*. While this result might sound limited due to the restricted class of simple PTGs we can handle, we recall that the previous works mentioned above [BLMR06, Rut11, HIJM13] have demonstrated that solving SPTGs is a key result towards solving more general PTGs. Moreover, this algorithm is, as far as we know, the first to handle the full class of SPTGs with arbitrary weights, and we note that the solutions (either the algorithms or the proofs) known so far do not generalise to this case. Notice also that previous algorithms provided exponential-time algorithms ($2^{O(n^2)}$ for [Rut11] and $O(12^n)$ for [HIJM13], with n the number of locations of the game), whereas we obtain a pseudo-polynomial time complexity (see Theorem 5.13 for the exact bound).¹ Finally, as a side result, this algorithm allows us to solve the more general class of *negative-reset-acyclic* one-clock PTGs that we introduce. This also improves the previous exponential complexity for one-clock PTGs with only non-negative prices to a pseudo-polynomial time complexity. However, the decidability (and thus complexity) of the whole class of one-clock PTGs with arbitrary weights remains open so far: our result may be seen as a potentially important milestone towards this goal.

2. QUANTITATIVE REACHABILITY GAMES

The semantics of the priced timed games we study in this work can be expressed in the setting of *quantitative reachability games* as defined below. Intuitively, in such a game, two players (Min and Max) play by changing alternatively the current configuration of the game. The game ends when it reaches a final configuration, and Min has to pay to Max a price associated with the sequence of configurations and of transitions taken (hence, Min is trying to minimise this price while Max wants to maximise it).

Note that the framework of quantitative reachability games that we develop here (and for which we prove a determinacy result, see Theorem 2.2) can be applied to other settings than priced timed games. For example, special cases of quantitative reachability games are

¹In the shorter version of this article, published in the proceedings of FSTTCS 2015 [BGH⁺15], only exponential-time complexity was provided: new techniques have allowed us to obtain the pseudo-polynomial time complexity.

finite quantitative reachability games—where the set of configurations is finite—that have been thoroughly studied in [BGHM16] under the name of *min-cost reachability games*. In this article, we will rely on quantitative reachability games with *uncountably* many states as the semantics of priced timed games. Similarly, our quantitative reachability games could be used to formalise the semantics of hybrid games [BBC06, BBJ⁺08] or any (non-probabilistic) game with a reachability objective. We start our discussion by defining formally those games:

Definition 2.1 (Quantitative reachability games). A *quantitative reachability game* is a tuple $G = (C_{\text{Min}}, C_{\text{Max}}, F, \Sigma, E, \text{Cost})$, where $C = C_{\text{Min}} \uplus C_{\text{Max}} \uplus F$ is the set of *configurations* (that does not need to be finite, nor countable), partitioned into the set C_{Min} of configurations of player **Min**, the set C_{Max} of configurations of player **Max**, and the set F of *final configurations*; Σ is a (potentially infinite) alphabet whose elements are called letters; $E \subseteq (C \setminus F) \times \Sigma \times C$ is the transition relation; and $\text{Cost}: (C\Sigma)^*C \rightarrow \mathbb{R}$ maps each finite sequence $c_1a_1 \cdots a_{n-1}c_n$ to a real number called the *cost* of $c_1a_1 \cdots a_{n-1}c_n$.

A *play* is a finite or infinite sequence $\rho = c_1a_1c_2 \cdots$ alternating between configurations and letters, and such that for all $i \geq 0$: $(c_i, a_i, c_{i+1}) \in E$. The length of a play ρ , denoted $|\rho|$ is the number of configuration occurring in it. As such, if ρ is infinite, $|\rho| = +\infty$. For the sake of clarity, we denote a play $c_1a_1c_2 \cdots$ as $c_1 \xrightarrow{a_1} c_2 \cdots$. A *completed play* is either (1) an infinite play, or (2) a finite play ending in a deadlock, i.e. a configuration c such that the set $\{(c, a, c') \in E \mid a \in \Sigma \wedge c' \in C\}$ is empty. Note that every play reaching a final state ends in a deadlock, hence infinite plays never visit F .

We take the viewpoint of player **Min** who wants to reach a final configuration. Thus, the *price* of a completed play $\rho = c_1 \xrightarrow{a_1} c_2 \cdots$, denoted $\text{Price}(\rho)$ is either $+\infty$ if either $|\rho| = +\infty$ or ρ is a finite play that does not end in a final state (this is the worst situation for **Min**, which explains why the price is maximal in this case); or $\text{Cost}(c_1 \xrightarrow{a_1} c_2 \cdots c_n)$ if $|\rho| = n$ and $c_n \in F$.

A *strategy* for player **Min** is a function σ_{Min} mapping all finite plays ending in a configuration $c \in C_{\text{Min}}$ (excluding deadlocks) to a transition $(c, a, c') \in E$. Strategies σ_{Max} of player **Max** are defined accordingly. We let $\text{Strat}_{\text{Min}}(G)$ and $\text{Strat}_{\text{Max}}(G)$ be the sets of strategies of **Min** and **Max**, respectively. A pair $(\sigma_{\text{Min}}, \sigma_{\text{Max}}) \in \text{Strat}_{\text{Min}}(\mathcal{G}) \times \text{Strat}_{\text{Max}}(\mathcal{G})$ is called a *profile of strategies*. Together with an initial configuration c_1 , it defines a unique completed play $\text{CPlay}(c_1, \sigma_{\text{Min}}, \sigma_{\text{Max}}) = c_1 \xrightarrow{a_1} c_2 \cdots$ such that for all $i \geq 0$: $(c_i, a_i, c_{i+1}) = \sigma_{\text{Min}}(c_1 \xrightarrow{a_1} c_2 \cdots c_i)$ if $c_i \in C_{\text{Min}}$; and $(c_i, a_i, c_{i+1}) = \sigma_{\text{Max}}(c_1 \xrightarrow{a_1} c_2 \cdots c_i)$ if $c_i \in C_{\text{Max}}$. We let $\text{Play}(\sigma_{\text{Min}})$ (resp. $\text{CPlay}(\sigma_{\text{Min}})$) be the set of plays (resp. completed plays) that conform with σ_{Min} . That is, $c_1 \xrightarrow{a_1} c_2 \cdots \in \text{Play}(\sigma_{\text{Min}})$ iff for all $i \geq 0$: $c_i \in C_{\text{Min}}$ implies $(c_i, a_i, c_{i+1}) = \sigma_{\text{Min}}(c_1 \xrightarrow{a_1} c_2 \cdots c_i)$. We let $\text{Play}(c_1, \sigma_{\text{Min}})$ (resp. $\text{CPlay}(c_1, \sigma_{\text{Min}})$) be the subset of plays from $\text{Play}(\sigma_{\text{Min}})$ (resp. $\text{CPlay}(\sigma_{\text{Min}})$) that start in c_1 . We define $\text{Play}(\sigma_{\text{Max}})$ and $\text{Play}(c_1, \sigma_{\text{Max}})$ as well as the completed variants accordingly. Given an initial configuration c_1 , the price of a strategy σ_{Min} of **Min** is:

$$\text{Price}(c_1, \sigma_{\text{Min}}) = \sup_{\rho \in \text{CPlay}(c_1, \sigma_{\text{Min}})} \text{Price}(\rho).$$

It matches the intuition to be the largest price that **Min** may pay while following strategy σ_{Min} . This definition is equal to $\sup_{\sigma_{\text{Max}}} \text{Price}(\text{CPlay}(c_1, \sigma_{\text{Min}}, \sigma_{\text{Max}}))$, which is intuitively the highest price that **Max** can force **Min** to pay if **Min** follows σ_{Min} . Similarly, given a strategy σ_{Max} of **Max**

Max, we define the price of σ_{Max} as

$$\text{Price}(c_1, \sigma_{\text{Max}}) = \inf_{\rho \in \text{CPlay}(c_1, \sigma_{\text{Max}})} \text{Price}(\rho) = \inf_{\sigma_{\text{Min}}} \text{Price}(\text{CPlay}(c_1, \sigma_{\text{Min}}, \sigma_{\text{Max}})).$$

It corresponds to the least price that Min can achieve once Max has fixed its strategy σ_{Max} .

From there, two different definitions of the value of a configuration c_1 arise, depending on which player chooses its strategy first. The *upper value* of c_1 , defined as

$$\overline{\text{Val}}(c_1) = \inf_{\sigma_{\text{Min}}} \sup_{\sigma_{\text{Max}}} \text{Price}(\text{CPlay}(c_1, \sigma_{\text{Min}}, \sigma_{\text{Max}}))$$

corresponds to the least price that Min can ensure when choosing its strategy *before* Max, while the *lower value*, defined as

$$\underline{\text{Val}}(c_1) = \sup_{\sigma_{\text{Max}}} \inf_{\sigma_{\text{Min}}} \text{Price}(\text{CPlay}(c_1, \sigma_{\text{Min}}, \sigma_{\text{Max}}))$$

corresponds to the least price that Min can ensure when choosing its strategy *after* Max. It is easy to see that $\underline{\text{Val}}(c_1) \leq \overline{\text{Val}}(c_1)$, which explains the chosen names. Indeed, if Min picks its strategy after Max, it has more information, and then can, in general, choose a better response.

In general, the order in which players choose their strategies can modify the outcome of the game. However, for quantitative reachability games, this makes no difference, and the value is the same whichever player picks its strategy first. This result, known as the *determinacy property*, is formalised here:

Theorem 2.2 (Determinacy of quantitative reachability games). *For all quantitative reachability games G and configurations c_1 , $\overline{\text{Val}}(c_1) = \underline{\text{Val}}(c_1)$.*

Proof. To establish this result, we rely on a general determinacy result of Donald Martin [Mar75]. This result concerns *qualitative* games (i.e. games where players either win or lose the game, and do not pay a price), called Gale-Stewart games. So, we first explain how to reduce a quantitative reachability game $G = (C_{\text{Min}}, C_{\text{Max}}, F, \Sigma, E, \text{Cost})$ to a family of such Gale-Stewart games $\text{Threshold}(G, r)$ parametrised by a threshold $r \in \mathbb{R}$.

The Gale-Stewart game $\text{Threshold}(G, r)$ is played on an infinite tree whose vertices are owned by either of the players. A play is then a maximal branch in this tree, built as follows: the player who owns the root of the tree first picks a successor of the root that becomes the current vertex. Then, the player who owns this vertex gets to choose a successor that becomes the current one, etc. The game ends when a leaf is reached, where the winner is declared thanks to a given set Win of winning leaves.

In our case, the vertices of $\text{Threshold}(G, r)$ are the finite plays $c_1 \xrightarrow{a_1} c_2 \cdots c_n$ of G starting from configuration c_1 . Such a vertex $v = c_1 \xrightarrow{a_1} c_2 \cdots c_n$ is owned by Min iff $c_n \in C_{\text{Min}}$; otherwise v belongs to Max. A vertex $v = c_1 \xrightarrow{a_1} c_2 \cdots c_n$ has successors only if $c_n \notin F$. In this case, the successors of v are all the vertices $v \xrightarrow{a} c$ such that $(c_n, a, c) \in E$. Finally, a leaf $c_1 \xrightarrow{a_1} c_2 \cdots c_n$ (thus, with $c_n \in F$) is winning for Min iff $\text{Cost}(c_1 \xrightarrow{a_1} c_2 \cdots c_n) \leq r$.

As a consequence, the set of winning plays in $\text{Threshold}(G, r)$ is:

$$\text{Win} = \bigcup_{v \in L \text{ s.t. } \text{Cost}(v) \leq r} \{\text{branch}(v)\}$$

where L is the set of leaves of $Threshold(G, r)$, and $branch(v)$ is the (unique) branch from c_1 to v . Then, we rewrite the definition of Win as:

$$Win = \bigcup_{v \in L \text{ s.t. } Cost(v) \leq r} Cone(v)$$

where $Cone(v)$ is the set of plays in $Threshold(G, r)$ that visit v . Indeed, when v is a leaf, the set $Cone(v)$ reduces to the singleton containing only $branch(v)$. Thus, the set of winning plays (for Min) is an open set, defined in the topology generated from the $Cone(v)$ sets, and we can apply [Mar75] to conclude that $Threshold(G, r)$ is a determined game for all quantitative reachability games G and all thresholds $r \in \mathbb{R}$ i.e. either Min or Max has a winning strategy from the root of the tree. Moreover, notice that Min wins the game $Threshold(G, r)$ iff it guarantees in G an upper value $\overline{Val}(c_1) \leq r$. Similarly, Max wins the game $Threshold(G, r)$ iff it guarantees in G a lower value $\underline{Val}(c_1) > r$.

We rely on this result to prove that $\underline{Val}(c_1) \geq \overline{Val}(c_1)$ in G (the other inequality holds by definition of $\underline{Val}(c_1)$ and $\overline{Val}(c_1)$). We consider two cases:

- (1) If $\overline{Val}(c_1) = -\infty$, then, since $\underline{Val}(c_1) \leq \overline{Val}(c_1)$, we have $\underline{Val}(c_1) = -\infty$ too.
- (2) If $\overline{Val}(c_1) > -\infty$, consider any real number r such that $r < \overline{Val}(c_1)$. Therefore, Min loses in the game $Threshold(G, r)$. By determinacy, Max wins in this game, i.e. $\underline{Val}(c_1) \geq r$. Therefore, $r < \overline{Val}(c_1)$ implies $r \leq \underline{Val}(c_1)$. Thus, we have shown that, for all r : $r < \overline{Val}(c_1)$ implies $r \leq \underline{Val}(c_1)$. This is equivalent to saying that for all r : either $r \geq \overline{Val}(c_1)$, or $r \leq \underline{Val}(c_1)$. This can happen only when $\overline{Val}(c_1) \leq \underline{Val}(c_1)$. \square

Now that we have showed that quantitative reachability games are determined, we can denote by Val the value of the game, defined as $Val = \overline{Val} = \underline{Val}$.

3. PRICED TIMED GAMES

We are now ready to formally introduce the core model of this article: priced timed games. We start by the formal definition, then study some properties of the value function of those games (Section 3.2). Next, we introduce the restricted class of *simple priced timed games* (Section 3.3) and close this section by discussing some special strategies (called *switching strategies*) that we will rely upon in our algorithms to solve priced timed games.

3.1. Notations and definitions. As usual, we let \mathbb{N} , \mathbb{Z} , \mathbb{Q} , \mathbb{R} , and \mathbb{R}^+ be the set of non-negative integers, integers, rational numbers, real numbers, and non-negative real numbers respectively. We also let $\overline{\mathbb{R}} = \mathbb{R} \cup \{+\infty, -\infty\}$. Let x denote a non-negative real-valued variable called *clock*. A *guard* (or *clock constraint*) is an interval with endpoints in $\mathbb{Q} \cup \{+\infty\}$. We often abbreviate guards, writing for instance $x \leq 5$ instead of $[0, 5]$. The set of all guards on the clock x is called $\text{Guard}(x)$. Let $S \subseteq \text{Guard}(x)$ be a finite set of guards. We let $\llbracket S \rrbracket = \bigcup_{I \in S} I$. Assuming $M_0 = 0 < M_1 < \dots < M_k$ are all the endpoints of the intervals in S (to which we add 0 if needed), we let

$$\text{Reg}_S = \{(M_i, M_{i+1}) \mid 0 \leq i \leq k-1\} \cup \{\{M_i\} \mid 0 \leq i \leq k\}$$

be the set of *regions* of S . Thus, intuitively, a region of S is either an open interval whose endpoints are consecutive endpoints the intervals in S , or singletons containing one endpoint from S . Observe that Reg_S is also a set of guards.

We rely on the notion of *cost function* to formalise the notion of optimal value function sketched in the introduction. Formally, for a set of guards $S \subseteq \text{Guard}(x)$, a *cost function* over S is a function $f: \llbracket \text{Reg}_S \rrbracket \rightarrow \overline{\mathbb{R}}$ such that over each region $r \in \text{Reg}_S$, f is either infinite or it is a continuous piecewise affine function with rational slopes and a finite set of rational cutpoints (points where the first derivative is not defined). In particular, if $f(r) = \{f(\nu) \mid \nu \in r\}$ contains $+\infty$ (respectively, $-\infty$) for some region r , then $f(r) = \{+\infty\}$ ($f(r) = \{-\infty\}$). We denote by CF_S the set of all cost functions over S .

In our algorithm to solve SPTGs, we will need to combine cost functions thanks to the \triangleright operator. Let $f \in \text{CF}_S$ and $f' \in \text{CF}_{S'}$ be two cost functions on sets of guards $S, S' \subseteq \text{Guard}(x)$, such that $\llbracket S \rrbracket \cap \llbracket S' \rrbracket$ is a singleton. We let $f \triangleright f'$ be the cost function in $\text{CF}_{S \cup S'}$ such that $(f \triangleright f')(\nu) = f(\nu)$ for all $\nu \in \llbracket \text{Reg}_S \rrbracket$, and $(f \triangleright f')(\nu) = f'(\nu)$ for all $\nu \in \llbracket \text{Reg}_{S'} \rrbracket \setminus \llbracket \text{Reg}_S \rrbracket$. For example, let $S = \{\{0\}, (0, 1), \{1\}\}$ and $S' = \{\{1\}\}$. We define the cost functions f_1 and f_2 such that f_1 is equal to $+\infty$ on the set of regions Reg_S and f_2 is equal to 0 on the set of regions $\text{Reg}_{S'}$. The cost function $f_2 \triangleright f_1 \in \text{CF}_{S \cup S'}$ is equal to $+\infty$ on $[0, 1)$ and to 0 on $\{1\}$ and the cost function $f_1 \triangleright f_2 \in \text{CF}_{S'}$ is equal to $+\infty$ on $[0, 1]$. Thus $f_1 \triangleright f_2$ is equal to f_1 while $f_2 \triangleright f_1$ extends f_2 with a $+\infty$ value on $[0, 1)$.

We consider an extended notion of one-clock priced timed games (PTGs for short) allowing for the use of *urgent locations*, where only a zero delay can be spent, and *final cost functions* which are associated with all final locations and incur an extra cost to be paid when ending the game in this location:

Definition 3.1. A *priced timed game* (PTG for short) \mathcal{G} is a tuple $(L_{\text{Min}}, L_{\text{Max}}, L_f, L_u, \varphi, \Delta, \pi)$ where:

- L_{Min} and L_{Max} are finite sets of *locations* belonging respectively to player Min and Max. L_f is a finite set of *final locations*. We assume that $L_{\text{Min}}, L_{\text{Max}}$ and L_f are disjoint and denote $L = L_{\text{Min}} \uplus L_{\text{Max}} \uplus L_f$ the set of all locations of the PTG;
- $L_u \subseteq L \setminus L_f$ is the set of *urgent locations*²;
- $\Delta \subseteq (L \setminus L_f) \times \text{Guard}(x) \times \{\top, \perp\} \times L$ is a finite set of *transitions*. We denote by $S_{\mathcal{G}} = \{I \mid \exists \ell, R, \ell' : (\ell, I, R, \ell') \in \Delta\}$ the set of all guards occurring on some transitions of the PTG;
- $\varphi = (\varphi_\ell)_{\ell \in L_f}$ associates to all locations $\ell \in L_f$ a *final cost function*, that is an affine³ cost function φ_ℓ with rational coefficients;
- $\pi: (L \setminus L_f) \cup \Delta \rightarrow \mathbb{Z}$ is a mapping associating an integer weight to all non-final locations and transitions.

Intuitively, a transition (ℓ, I, R, ℓ') changes the current location from ℓ to ℓ' if the clock has value in I and the clock is reset according to the Boolean R . We assume that, in all PTGs, the clock x is *bounded*, i.e. there is $M \in \mathbb{N}$ such that for all guards $I \in S_{\mathcal{G}}$, $I \subseteq [0, M]$.⁴ We denote by $\text{Reg}_{\mathcal{G}}$ the set $\text{Reg}_{S_{\mathcal{G}}}$ of *regions of \mathcal{G}* . We further denote⁵ by $\Pi_{\mathcal{G}}^{\text{tr}}$, $\Pi_{\mathcal{G}}^{\text{loc}}$ and $\Pi_{\mathcal{G}}^{\text{fin}}$ respectively the values $\max_{\delta \in \Delta} |\pi(\delta)|$, $\max_{\ell \in (L \setminus L_f)} |\pi(\ell)|$ and $\sup_{\nu \in [0, M]} \max_{\ell \in L_f} |\varphi_\ell(\nu)| = \max_{\ell \in L_f} \max(|\varphi_\ell(0)|, |\varphi_\ell(M)|)$ (the last equality holds because

²Here we differ from [BLMR06] where $L_u \subseteq L_{\text{Max}}$.

³In our one-clock setting, an affine function is of the form $\varphi_\ell(\nu) = a \times \nu + b$.

⁴This last restriction is *not* without loss of generality in the case of PTGs. While all timed automata \mathcal{A} can be turned into an equivalent (with respect to reachability properties) \mathcal{A}' whose clocks are bounded [BFH⁺01], this technique cannot be applied to PTGs, in particular with arbitrary weights.

⁵Throughout the paper, we often drop the \mathcal{G} in the subscript of several notations when the game is clear from the context.

we have assumed that φ_ℓ is affine). That is, $\Pi_{\mathcal{G}}^{\text{tr}}$, $\Pi_{\mathcal{G}}^{\text{loc}}$ and $\Pi_{\mathcal{G}}^{\text{fin}}$ are the largest absolute values of the transition weights, location weights and final cost functions.

As announced in the first section, the semantics of a PTG $\mathcal{G} = (L_{\text{Min}}, L_{\text{Max}}, L_f, L_u, \varphi, \Delta, \pi)$ is given by a quantitative reachability game

$$G_{\mathcal{G}} = (C_{\text{Min}} = (L_{\text{Min}} \times \mathbb{R}^+), C_{\text{Max}} = (L_{\text{Max}} \times \mathbb{R}^+), F = (L_f \times \mathbb{R}^+), \Sigma = (\mathbb{R}^+ \times \Delta \times \mathbb{R}), E, \text{Cost})$$

that we describe now. Note that, from now on, we often confuse the PTG \mathcal{G} with its semantics $G_{\mathcal{G}}$, writing, for instance ‘the configurations of \mathcal{G} ’ instead of: ‘the configurations of $G_{\mathcal{G}}$ ’. We also lift the functions Price , $\overline{\text{Val}}$, $\underline{\text{Val}}$ and Val , and the notions of plays from $G_{\mathcal{G}}$ to \mathcal{G} . A *configuration* of \mathcal{G} is a pair $s = (\ell, \nu) \in L \times \mathbb{R}^+$, where ℓ and ν are respectively the current location and clock value of \mathcal{G} . We denote by $\text{Conf}_{\mathcal{G}}$ the set of all configurations of \mathcal{G} . Let (ℓ, ν) and (ℓ', ν') be two configurations, let $\delta = (\ell, I, R, \ell') \in \Delta$ be a transition of \mathcal{G} and $t \in \mathbb{R}^+$ be a delay. Then, $((\ell, \nu), (t, \delta, c), (\ell', \nu')) \in E$, iff:

- (1) $\ell \in L_u$ implies $t = 0$ (no time can elapse in urgent locations);
- (2) $\nu + t \in I$ (the guard is satisfied);
- (3) $R = \top$ implies $\nu' = 0$ (when the clock is reset);
- (4) $R = \perp$ implies $\nu' = \nu + t$ (when the clock is not reset);
- (5) $c = \pi(\delta) + t \times \pi(\ell)$ (the cost of (t, δ) takes into account the weight of ℓ , the delay t spent in ℓ , and the weight of δ).

In this case, we say that there is a (t, δ) -transition from (ℓ, ν) to (ℓ', ν') with cost c , and we denote this by $(\ell, \nu) \xrightarrow{t, \delta, c} (\ell', \nu')$. For two configurations s and s' , we also write $s \xrightarrow{c} s'$ whenever there are t and δ such that $s \xrightarrow{t, \delta, c} s'$. Observe that, since the alphabet of $G_{\mathcal{G}}$ is $\mathbb{R}^+ \times \Delta \times \mathbb{R}$, and its set of configurations is $\text{Conf}_{\mathcal{G}}$, plays of \mathcal{G} are of the form $\rho = (\ell_1, \nu_1) \xrightarrow{t_1, \delta_1, c_1} (\ell_2, \nu_2) \cdots$. Finally, the cost function Cost is obtained by summing the costs of the play (transitions and time spent in the locations) and the final cost function if applicable. Formally, let $\rho = (\ell_1, \nu_1) \xrightarrow{t_1, \delta_1, c_1} (\ell_2, \nu_2) \cdots (\ell_n, \nu_n)$ be a finite play such that for all $k < n$, $\ell_k \notin L_f$. Then, $\text{Cost}(\rho) = \sum_{i=1}^{n-1} c_i + \varphi_{\ell_n}(\nu_n)$ if $\ell_n \in L_f$, and $\text{Cost}(\rho) = \sum_{i=1}^{n-1} c_i$ otherwise.

As sketched in the introduction, we consider optimal reachability-price games on PTGs, where the aim of player Min is to reach a location of L_f while minimising the price. Since the semantics of PTGs is defined in terms of quantitative reachability games, we can apply Theorem 2.2, and deduce that all PTGs \mathcal{G} are determined. Hence, for all PTGs the value function Val is well-defined, and we denote it by $\text{Val}_{\mathcal{G}}$ when we need to emphasise the game it refers to.

For example, consider the PTG on the left of Figure 1. Using the final cost function φ constantly equal to 0, its value function for location ℓ_1 is represented on the right. The completed play $\rho = (\ell_1, 0) \xrightarrow{0, t_{1,2}, 0} (\ell_2, 0) \xrightarrow{1/4, t_{2,3}, -3.5} (\ell_3, 1/4) \xrightarrow{0, t_{3,7}, 6} (\ell_7, 1/4) \xrightarrow{3/4, t_{7,f}, -12} (\ell_f, 1)$ where $t_{n,m} = (\ell_n, [0, 1], \perp, \ell_m)$ ends in the unique final location ℓ_f and its price is $\text{Price}(\rho) = 0 - 3.5 + 6 - 12 = -9.5$.

Let us fix a PTG \mathcal{G} with initial configuration c_1 . We say that a strategy σ_{Min} of Min is *optimal* if $\text{Price}(c_1, \sigma_{\text{Min}}) = \text{Val}_{\mathcal{G}}(c_1)$, i.e. it ensures Min to enforce the value of the game, whatever Max does. Similarly, σ_{Min} is ε -*optimal*, for $\varepsilon > 0$, if $\text{Price}(c_1, \sigma_{\text{Min}}) \leq \text{Val}_{\mathcal{G}}(c_1) + \varepsilon$. And, symmetrically, a strategy σ_{Max} of Max is *optimal* (respectively, ε -*optimal*) if $\text{Price}(c_1, \sigma_{\text{Max}}) = \text{Val}_{\mathcal{G}}(c_1)$ (respectively, $\text{Price}(c_1, \sigma_{\text{Max}}) \geq \text{Val}_{\mathcal{G}}(c_1) - \varepsilon$).

3.2. Properties of the value. Let us now discuss useful preliminary properties of the value functions of PTGs. We have already shown the determinacy of the game, ensuring the existence of the value function. We will now establish a stronger (and, to the best of our knowledge, original) result. For all locations ℓ , let $\text{Val}_{\mathcal{G}}(\ell)$ denote the function such that $\text{Val}_{\mathcal{G}}(\ell)(\nu) = \text{Val}_{\mathcal{G}}(\ell, \nu)$ for all $\nu \in \mathbb{R}^+$. Then, we show that, for all ℓ , $\text{Val}_{\mathcal{G}}(\ell)$ is a *piecewise continuous function* that might exhibit discontinuities *only on the borders of the regions* of $\text{Reg}_{\mathcal{G}}$.

Theorem 3.2. *For all (one-clock) PTGs \mathcal{G} , for all $r \in \text{Reg}_{\mathcal{G}}$, for all $\ell \in L$, $\text{Val}_{\mathcal{G}}(\ell)$ is either infinite or continuous over r .*

Proof. The main ingredient of our proof is, given a strategy σ_{Min} of Min , a location ℓ of the game, a region $r \in \text{Reg}_{\mathcal{G}}$ and valuations $\nu, \nu' \in r$, to show how to build a strategy σ'_{Min} and a function g such that g maps plays starting in (ℓ, ν') and consistent with σ'_{Min} to plays starting in (ℓ, ν) consistent with σ_{Min} with similar behaviour and cost. More precisely, we define σ'_{Min} and g by induction on the length of the finite play that is given as argument and rely on the following set of induction hypothesis:

Induction hypothesis: There exists a strategy σ'_{Min} and a function g from the plays of length (respectively) $k-1$ and k , starting in (ℓ, ν') and consistent with σ'_{Min} , to (respectively) transitions in \mathcal{G} and plays starting in (ℓ, ν) , consistent with σ_{Min} , such that for all plays

$\rho' = (\ell_1, \nu'_1) \xrightarrow{c'_1} \dots \xrightarrow{c'_{k-1}} (\ell_k, \nu'_k)$ starting in (ℓ, ν') and consistent with σ'_{Min} , denoting $(\ell_1, \nu_1) \xrightarrow{c_1} \dots \xrightarrow{c_{k'-1}} (\ell_{k'}, \nu_{k'})$ the play $g(\rho')$ starting in (ℓ, ν) we have:

- (1) ρ' and $g(\rho')$ have the same length, i.e. $k' = k$,
- (2) for every $i \in \{1, \dots, k\}$, ν_i and ν'_i are in the same region, i.e. there exists a region $r' \in \text{Reg}_{\mathcal{G}}$ such that $\nu_i \in r'$ and $\nu'_i \in r'$,
- (3) $|\nu_k - \nu'_k| \leq |\nu - \nu'|$,
- (4) $\text{Cost}(\rho') \leq \text{Cost}(g(\rho')) + \Pi^{\text{loc}}(|\nu - \nu'| - |\nu_k - \nu'_k|)$.

Notice that no property is required on the strategy σ'_{Min} for finite plays that do not start in (ℓ, ν') .

Let us explain how this result would imply the theorem before going through the induction itself. Let $r \in \text{Reg}_{\mathcal{G}}$ be a region of the game and ℓ be a location. Remark first that the result directly implies that if the value of the game is finite for some valuation ν in r , then it is finite for all other valuation ν' in r . Indeed, a finite value of the game in (ℓ, ν) implies that there exists a strategy σ_{Min} such that every play consistent with it and starting in (ℓ, ν) reaches a final location with a time valuation such that the final cost function is finite. Moreover, denoting σ'_{Min} the strategy obtained from σ_{Min} thanks to the above result, any play ρ' starting in (ℓ, ν') and consistent with σ'_{Min} reaches a final location (since $g(\rho')$ does) and the final cost function is finite as the final time valuation of ρ' and $g(\rho')$ sit in the same region and, by definition, a final cost function is either always finite or always infinite within a region.

Now, assuming the value of the game is finite over r , in order to show that $\text{Val}_{\mathcal{G}}(\ell)$ is continuous over r , we need to show that, for all $\nu \in r$, for all $\varepsilon > 0$, there exists $\delta > 0$ s.t. for all $\nu' \in r$ with $|\nu - \nu'| \leq \delta$, we have $|\text{Val}(\ell, \nu) - \text{Val}(\ell, \nu')| \leq \varepsilon$. To this end, we can show that:

$$|\text{Val}(\ell, \nu) - \text{Val}(\ell, \nu')| \leq (\Pi^{\text{loc}} + K_{\text{fin}})|\nu - \nu'| \quad (3.1)$$




Figure 2: The definition of t' when (a) $\nu_k' \leq \nu_k$, (b) $\nu_k < \nu_k' < \nu_k + t$, (c) $\nu_k < \nu_k + t < \nu_k'$.

where K_{fin} is the greatest absolute value of the slopes appearing in the piecewise affine functions within φ . Indeed, assume that this inequality holds, and consider a clock value $\nu \in r$ and a positive real number ε . Then, we let $\delta = \frac{\varepsilon}{\Pi^{loc} + K_{fin}}$, and we consider a valuation ν' s.t. $|\nu - \nu'| \leq \delta$. In this case, equation (3.1) becomes:

$$|\text{Val}(\ell, \nu) - \text{Val}(\ell, \nu')| \leq (\Pi^{loc} + K_{fin})|\nu - \nu'| \leq (\Pi^{loc} + K_{fin}) \frac{\varepsilon}{\Pi^{loc} + K_{fin}} \leq \varepsilon.$$

Thus, proving equation (3.1) is sufficient to establish continuity. On the other hand, equation (3.1) is equivalent to:

$$\text{Val}(\ell, \nu) \leq \text{Val}(\ell, \nu') + (\Pi^{loc} + K_{fin})|\nu - \nu'| \quad \text{and} \quad \text{Val}(\ell, \nu') \leq \text{Val}(\ell, \nu) + (\Pi^{loc} + K_{fin})|\nu - \nu'|.$$

As those two last equations are symmetric with respect to ν and ν' , we only have to show either of them. We thus focus on the latter, which, by using the upper value, can be reformulated as: for all strategies σ_{Min} of Min, there exists a strategy σ'_{Min} such that $\text{Price}((\ell, \nu'), \sigma'_{\text{Min}}) \leq \text{Price}((\ell, \nu), \sigma_{\text{Min}}) + (\Pi^{loc} + K_{fin})|\nu - \nu'|$. Note that this last equation is equivalent to say that there exists a function g mapping plays ρ' from (ℓ, ν') , consistent with σ'_{Min} (i.e. such that $\rho' = \text{Play}((\ell, \nu'), \sigma'_{\text{Min}}, \sigma_{\text{Max}})$ for some strategy σ_{Max} of Max) to plays from (ℓ, ν) , consistent with σ_{Min} , such that, for all such ρ' the final time valuations of ρ' and $g(\rho')$ differ by at most $|\nu - \nu'|$ and:

$$\text{Cost}(\rho') \leq \text{Cost}(g(\rho')) + \Pi^{loc}|\nu - \nu'|$$

which is exactly what our claimed induction achieves. Thus, to conclude this proof, let us now define σ'_{Min} and g , by induction on the length k of ρ' .

Base case $k = 1$: In this case, σ'_{Min} does not have to be defined. Moreover, in that case, $\rho' = (\ell, \nu')$ and $g(\rho') = (\ell, \nu)$. Both plays have length 1, ν and ν' are in the same region by hypothesis, and $\text{Cost}(\rho') = \text{Cost}(g(\rho')) = 0$, therefore all four properties are true.

Inductive case: Let us suppose now that the construction is done for a given $k \geq 1$, and perform it for $k + 1$. We start with the construction of σ'_{Min} . To that extent, consider

a play $\rho' = (\ell_1, \nu'_1) \xrightarrow{c'_1} \dots \xrightarrow{c'_{k-1}} (\ell_k, \nu'_k)$ from (ℓ, ν') , consistent with σ'_{Min} (provided by induction hypothesis) such that ℓ_k is a location of player Min. Let t and δ be the choice of delay and transition made by σ_{Min} on $g(\rho')$, i.e. $\sigma_{\text{Min}}(g(\rho')) = (t, \delta)$. Then, we define $\sigma'_{\text{Min}}(\rho') = (t', \delta)$ where $t' = \max(0, \nu_k + t - \nu'_k)$. The delay t' respects the guard of transition δ , as can be seen from Figure 2. Indeed, either $\nu_k + t = \nu'_k + t'$ (cases (a) and (b) in the figure) or $\nu_k \leq \nu_k + t \leq \nu'_k$, in which case ν'_k is in the same region as $\nu_k + t$ since ν_k and ν'_k are in the same region by induction hypothesis.

Let us now build the mapping g . Let $\rho' = (\ell_1, \nu'_1) \xrightarrow{c'_1} \dots \xrightarrow{c'_k} (\ell_{k+1}, \nu'_{k+1})$ be a play from (ℓ, ν') consistent with σ'_{Min} and let $\tilde{\rho}' = (\ell_1, \nu'_1) \xrightarrow{c'_1} \dots \xrightarrow{c'_{k-1}} (\ell_k, \nu'_k)$ its prefix of length k . Let (t', δ) be the delay and transition taken after $\tilde{\rho}'$. Using the construction of g over plays of length k by induction, the play $g(\tilde{\rho}') = (\ell_1, \nu_1) \xrightarrow{c_1} \dots \xrightarrow{c_{k-1}} (\ell_k, \nu_k)$ (with $(\ell_1, \nu_1) = (\ell, \nu)$) verifies properties (1), (2), (3) and (4). Then:

- if ℓ_k is a location of Min and $\sigma_{\text{Min}}(g(\tilde{\rho}')) = (t, \delta)$, then $g(\rho') = g(\tilde{\rho}') \xrightarrow{c_k} (\ell_{k+1}, \nu_{k+1})$ is obtained by applying those choices on $g(\tilde{\rho}')$;
- if ℓ_k is a location of Max , the last clock value ν_{k+1} of $g(\rho')$ is rather obtained by choosing action (t, δ) verifying $t = \max(0, \nu'_k + t' - \nu_k)$. Note that transition δ is allowed since both $\nu_k + t$ and $\nu'_k + t'$ are in the same region (for similar reasons as above).

By induction hypothesis $|\tilde{\rho}'| = |g(\tilde{\rho}')|$, thus: 1 holds, i.e. $|\rho'| = |g(\rho')|$. Moreover, ν_{k+1} and ν'_{k+1} are also in the same region as either they are equal to $\nu_k + t$ and $\nu'_k + t'$, respectively, or δ contains a reset in which case $\nu_{k+1} = \nu'_{k+1} = 0$ which proves 2. To prove 3, notice that we always have either $\nu_k + t = \nu'_k + t'$ or $\nu_k \leq \nu_k + t \leq \nu'_k = \nu'_k + t'$ or $\nu'_k \leq \nu'_k + t \leq \nu_k = \nu_k + t$. In all of these possibilities, we have $|(\nu_k + t) - (\nu'_k + t')| \leq |\nu_k - \nu'_k|$. We finally check property 4. In both cases:

$$\begin{aligned} \text{Cost}(\rho') &= \text{Cost}(\tilde{\rho}') + \pi(\delta) + t'\pi(\ell_k) \\ &\leq \text{Cost}(g(\tilde{\rho}')) + \Pi^{\text{loc}}(|\nu - \nu'| - |\nu_k - \nu'_k|) + \pi(\delta) + t'\pi(\ell_k) \\ &= \text{Cost}(g(\rho')) + (t' - t)\pi(\ell_k) + \Pi^{\text{loc}}(|\nu - \nu'| - |\nu_k - \nu'_k|). \end{aligned}$$

Let us prove that

$$|t' - t| \leq |\nu_k - \nu'_k| - |\nu'_{k+1} - \nu_{k+1}|. \quad (3.2)$$

If δ contains no reset, $t' = \nu'_{k+1} - \nu'_k$ and $t = \nu_{k+1} - \nu_k$, we have $|t' - t| = |\nu'_{k+1} - \nu'_k - (\nu_{k+1} - \nu_k)|$. Then, two cases are possible: either $t' = \max(0, \nu_k + t - \nu'_k)$ or $t = \max(0, \nu'_k + t' - \nu_k)$. So we have three different possibilities:

- if $t' + \nu'_k = t + \nu_k$, then $\nu'_{k+1} = \nu_{k+1}$, thus $|t' - t| = |\nu_k - \nu'_k| = |\nu_k - \nu'_k| - |\nu'_{k+1} - \nu_{k+1}|$.
- if $t = 0$, then $\nu_k = \nu_{k+1} \geq \nu'_{k+1} \geq \nu'_k$, thus $|\nu'_{k+1} - \nu'_k - (\nu_{k+1} - \nu_k)| = \nu'_{k+1} - \nu'_k = (\nu_k - \nu'_k) - (\nu_{k+1} - \nu'_{k+1}) = |\nu_k - \nu'_k| - |\nu'_{k+1} - \nu_{k+1}|$.
- if $t' = 0$, then $\nu'_k = \nu'_{k+1} \geq \nu_{k+1} \geq \nu_k$, thus $|\nu'_{k+1} - \nu'_k - (\nu_{k+1} - \nu_k)| = \nu_{k+1} - \nu_k = (\nu'_k - \nu_k) - (\nu'_{k+1} - \nu_{k+1}) = |\nu_k - \nu'_k| - |\nu'_{k+1} - \nu_{k+1}|$.

If δ contains a reset, then $\nu'_{k+1} = \nu_{k+1}$. If $t' = \nu_k + t - \nu'_k$, we have that $|t' - t| = |\nu_k - \nu'_k|$. Otherwise, $t' = 0$ and $t \leq \nu'_k - \nu_k$. In all cases, we have proved (3.2).

Together with the fact that $|\pi(\ell_k)| \leq \Pi^{\text{loc}}$, we conclude that:

$$\text{Cost}(\rho') \leq \text{Cost}(g(\rho')) + \Pi^{\text{loc}}(|\nu - \nu'| - |\nu_{k+1} - \nu'_{k+1}|).$$

Now that σ'_{Min} and g are defined (noticing that g is stable by prefix, we extend naturally its definition to infinite plays), notice that for all plays ρ' from (ℓ, ν') consistent with σ'_{Min} , either ρ' does not reach a final location and its price is $+\infty$, but in this case $g(\rho')$ has also price $+\infty$; or ρ' is finite. In this case, let ν'_k be the clock value of its last configuration, and ν_k be the clock value of the last configuration of $g(\rho')$. Combining (3) and (4) we have $\text{Cost}(\rho') \leq \text{Cost}(g(\rho')) + \Pi^{\text{loc}}|\nu - \nu'|$ which concludes the induction. \square

Remark 3.3. Let us consider the example in Figure 3 (that we describe informally since we did not properly define games with multiple clocks), with clocks x and y . One can easily check




Figure 3: A PTG with 2 clocks whose value function is not continuous inside a region.

that, starting from a configuration $(\ell_0, 0, 0.5)$ in location ℓ_0 and where $x = 0$ and $y = 0.5$, the following cycle can be taken: $(\ell_0, 0, 0.5) \xrightarrow{0, \delta_0, 0} (\ell_1, 0, 0.5) \xrightarrow{0.5, \delta_1, 2.5} (\ell_2, 0.5, 0) \xrightarrow{0.5, \delta_2, -2.5} (\ell_0, 0, 0.5)$, where δ_0 , δ_1 and δ_2 denote respectively the transitions from ℓ_0 to ℓ_1 ; from ℓ_1 to ℓ_2 ; and from ℓ_2 to ℓ_0 . Observe that the cost of this cycle is null, and that no other delays can be played, hence $\overline{\text{Val}}(\ell_0, 0, 0.5) = 0$. However, starting from a configuration $(\ell_0, 0, 0.6)$, and following the same path, yields the cycle $(\ell_0, 0, 0.6) \xrightarrow{0, \delta_0, 0} (\ell_1, 0, 0.6) \xrightarrow{0.4, \delta_1, 2} (\ell_2, 0.4, 0) \xrightarrow{0.6, \delta_2, -3} (\ell_0, 0, 0.6)$ with cost -1 . Hence, $\overline{\text{Val}}(\ell_0, 0, 0.6) = -\infty$, and the function is not continuous although both clocks values $(0, 0.5)$ and $(0, 0.6)$ are in the same region. Observe that this holds even for priced timed *automata*, since our example requires only one player.

3.3. Simple priced timed games. As sketched in the introduction, our main contribution is to solve the special case of simple one-clock priced timed games with arbitrary weights, where the clock is never reset and takes values in some fixed interval $[0, r]$ only. Formally, an r -SPTG, with $r \in \mathbb{Q}^+ \cap [0, 1]$, is a PTG $\mathcal{G} = (L_{\text{Min}}, L_{\text{Max}}, L_f, L_u, \varphi, \Delta, \pi)$ such that for all transitions $(\ell, I, R, \ell') \in \Delta$, $I = [0, r]$ (the clock is also bounded by r) and $R = \perp$. Hence, transitions of r -SPTGs are henceforth denoted by (ℓ, ℓ') , dropping the guard and the reset. Then, an SPTG is a 1-SPTG. This paper is mainly devoted to prove the following result on SPTGs.

Theorem 3.4. *Let \mathcal{G} be an SPTG. Then, for all locations $\ell \in L$, either $\text{Val}(\ell) \in \{-\infty, +\infty\}$, or $\text{Val}(\ell)$ is continuous and piecewise-affine with at most a pseudo-polynomial number of cutpoints (in the size of \mathcal{G}), i.e. polynomial if the prices of the game are encoded in unary. The value functions $\text{Val}(\ell)$ for all locations ℓ , as well as a pair of optimal strategies $(\sigma_{\text{Min}}, \sigma_{\text{Max}})$ (that always exist when no values are infinite) can be computed in pseudo-polynomial time.*

3.3.1. Proof strategy. Let us now highlight the main steps that will allow us to establish this theorem. The central argument consists in showing that all SPTGs admit ‘well-behaved’ optimal strategies for both players, in the sense that these strategies can be finitely described (and computed in pseudo-polynomial time). To this end, we rely on several new definitions that we are about to introduce and that we first describe informally.

We start by the case of **Max**: we will show that **Max** always has a *positional* (aka memoryless) optimal strategy. However, this is not sufficient to show that **Max** has an optimal strategy that can be *finitely* described: indeed, in the case of SPTGs, a positional




Figure 4: An SPTG where Min needs memory to play optimally

strategy associates a move to each *configuration* of the game, and there are uncountably many such configurations because of the possible values of the clock. Thus, we introduce the notion of *finite positional strategies* (FP-strategies for short). Such strategies partition the set $[0, 1]$ of possible clock values into finitely many intervals, and ensure that the same move is played throughout each interval: this move can be either to *wait* until the clock reaches the end of the interval, or to *take immediately* a given transition.

The case of Min is more involved, as shown in the next example taken from [BGHM16]. Consider the SPTG of Figure 4, where W is a positive integer, and every location has weight 0 (thus, it is an *untimed* game, as originally studied). We claim that the values of locations l_1 and l_2 are both $-W$. Indeed, consider the following strategy for Min: during each of the first W visits to l_2 (if any), go to l_1 ; else, go to l_f . Clearly, this strategy ensures that the final location l_f will eventually be reached, and that:

- (1) either transition (l_1, l_f) (with weight $-W$) will eventually be traversed;
- (2) or transition (l_1, l_2) (with weight -1) will be traversed at least W times.

Hence, in all plays following this strategy, the price will be at most $-W$. This strategy allows Min to secure $-W$, but it cannot ensure a lower price, since Max always has the opportunity to take the transition (l_1, l_f) (with cost $-W$) instead of cycling between l_1 and l_2 . Hence, Max's optimal choice is to follow the transition (l_1, l_f) as soon as l_1 is reached, securing a price of $-W$. The strategy we have just given is optimal for Min, and there are *no optimal memoryless strategies* for Min. Indeed, always playing (l_2, l_f) does not ensure a price at most $-W$; and, always playing (l_2, l_1) does not guarantee to reach the target, and this strategy has thus value $+\infty$.

This example shows the kind of strategies that we will prove sufficient for Min to play optimally: *first* play an FP-strategy to obtain a play prefix with a sufficiently low cost, by forcing negative cycles (if any); *second* play another FP-strategy that ensures that the target will eventually be reached. Such strategies have been introduced as *switching strategies* in [BGHM16], and can be finitely described by a pair $(\sigma_{\text{Min}}^1, \sigma_{\text{Min}}^2)$ of FP-strategies and a threshold K to trigger the switch when the length of the play prefix is at least K .

Computing the latter of these two strategies is easy: σ_{Min}^2 is basically an *attractor* strategy, which guarantees Min to reach the target (when possible) at a bounded cost. Thus, the main difficulty in identifying *optimal switching strategies* is to characterise σ_{Min}^1 . To do so, we further introduce the notion of *negative cycle strategies* (NC-strategies for short). Those strategies are FP-strategies which guarantee that all cycles taken have cost of -1 at most, without necessarily guaranteeing to eventually reach the target (as this will be taken care of by σ_{Min}^2). Among those NC-strategies, we identify so-called *fake optimal* strategies. Those are the NC-strategies that guarantee Min to obtain the optimal value (or better) when

the target is reached, but do not necessarily guarantee to reach the target (they are thus not *really* optimal, hence the name fake-optimal).

Based on these definitions, we will show that *all SPTGs* with only finite values admit such optimal switching strategies for Min and optimal FP-strategies for Max. By definition, these strategies can be finitely described (as a matter of fact, we will show that we can compute them in pseudo-polynomial time). Let us now give the formal definitions of those notions.

3.3.2. Finite positional strategies. We start with the notion of finite positional strategies, that will formalise a class of optimal strategies for Max:

Definition 3.5 (FP-strategies). A strategy σ is a *finite positional strategy* (FP-strategy for short) iff it is a memoryless strategy (i.e. for all finite plays $\rho_1 = \rho'_1 \xrightarrow{c_1} s$ and $\rho_2 = \rho'_2 \xrightarrow{c_2} s$ ending in the same configuration, we have $\sigma(\rho_1) = \sigma(\rho_2)$) and for all locations ℓ , there exists a finite sequence of rationals $0 = \nu_0^\ell < \nu_1^\ell < \nu_2^\ell < \dots < \nu_k^\ell = 1$ and a finite sequence of transitions $\delta_0, \dots, \delta_{2k} \in \Delta$ such that

- (1) for all $1 \leq i \leq k$ and $\nu \in (\nu_{i-1}^\ell, \nu_i^\ell)$, either $\sigma(\ell, \nu) = (0, \delta_{2i-1})$, or $\sigma(\ell, \nu) = (\nu_i^\ell - \nu, \delta_{2i-1})$;
- (2) for all $0 \leq i \leq k-1$, either $\sigma(\ell, \nu_i^\ell) = (0, \delta_{2i})$, or $\sigma(\ell, \nu_i^\ell) = (\nu_{i+1}^\ell - \nu_i^\ell, \delta_{2i})$; and
- (3) $\sigma(\ell, \nu_k^\ell) = (0, \delta_{2k})$.

We let $\text{pts}(\sigma)$ be the set of ν_i^ℓ for all ℓ and i , and $\text{int}(\sigma)$ be the set of all successive open intervals and singletons generated by $\text{pts}(\sigma)$. Finally, we let $|\sigma| = |\text{int}(\sigma)|$ be the size of σ . Intuitively, in each location ℓ and interval $(\nu_{i-1}^\ell, \nu_i^\ell)$, σ always returns the same move: either to take *immediately* δ_{2i-1} or to wait until the clock reaches the endpoint ν_i^ℓ and then take δ_{2i-1} (point 1 of the definition above). A similar behaviour also happens on the endpoints (point 2).

3.3.3. Switching strategies. On top of the definition of FP-strategies, we can now define the notion of switching strategy:

Definition 3.6 (Switching strategies). A switching strategy is described by a pair $(\sigma_{\text{Min}}^1, \sigma_{\text{Min}}^2)$ of FP-strategies and a switch threshold K . It consists in playing σ_{Min}^1 until the play contains K transitions (i.e. the length of the play prefix is at least $K+1$); and then to *switch* to strategy σ_{Min}^2 .

The role of σ_{Min}^2 is to ensure reaching a final location: it is thus a (classical) attractor strategy. The role of σ_{Min}^1 , on the other hand, is to allow Min to decrease the cost low enough (possibly by forcing negative cycles) to secure a price sufficiently low, and the computation of σ_{Min}^1 is thus the critical point in the computation of an optimal switching strategy. In the SPTG of Figure 4, for example, σ_{Min}^1 is the strategy that goes from ℓ_2 to ℓ_1 , σ_{Min}^2 is the strategy going directly to ℓ_f and the switch occurs after the threshold of $K = 2W$. The value of the game under this strategy is thus $-W$.

3.3.4. *Negative cycle strategies.* To characterise σ_{Min}^1 , we introduce now the notion of negative cycle strategy (NC-strategy):

Definition 3.7 (Negative cycle strategies). An NC-strategy σ_{Min} of Min is an FP-strategy such that for all plays $\rho = (\ell_1, \nu) \xrightarrow{c_1} \dots \xrightarrow{c_{k-1}} (\ell_k, \nu') \in \text{Play}(\sigma_{\text{Min}})$ with $\ell_1 = \ell_k$, and ν, ν' in the same interval of $\text{int}(\sigma_{\text{Min}})$, the sum of weights of *discrete transitions* is at most -1 , i.e. $\pi(\ell_1, \ell_2) + \dots + \pi(\ell_{k-1}, \ell_k) \leq -1$.

Let us now show that this definition allows one to find an upper bound on the costs of the plays following such an NC-strategy σ_{Min} .

Lemma 3.8. *Let σ_{Min} be an NC-strategy, and let $\rho \in \text{Play}(\sigma_{\text{Min}})$ be a finite play. Then:*

$$\text{Cost}(\rho) \leq \Pi^{\text{loc}} + (2|\sigma_{\text{Min}}| - 1) \times |L|\Pi^{\text{tr}} - \frac{|\rho|}{|L|} + 2|\sigma_{\text{Min}}|.$$

Proof. We start by proving a bound on the cost of a finite play $\tilde{\rho} \in \text{Play}(\sigma_{\text{Min}})$ such that all clock values are in the same interval I of $\text{int}(\sigma_{\text{Min}})$ and ending in (ℓ_f, ν_f) . In this case, we claim that:

$$\text{Cost}(\tilde{\rho}) \leq |I|\Pi^{\text{loc}} + |L|\Pi^{\text{tr}} - \lfloor (|\tilde{\rho}| - 1)/|L| \rfloor \quad (3.3)$$

(where $\lfloor \cdot \rfloor$ is the floor function). Indeed, the cost of $\tilde{\rho}$ is the sum of the weights generated while spending time in locations, plus the discrete weights of taking transitions. The former is bounded by $|I|\Pi^{\text{loc}}$ since the total time spent in locations is bounded by $|I|$. For the discrete weights, one can delete from $\tilde{\rho}$ at least $\lfloor (|\tilde{\rho}| - 1)/|L| \rfloor$ cycles, which all have a weight bounded above by -1 , since σ_{Min} is an NC-strategy. After removing those cycles from $\tilde{\rho}$, one ends up with a play of length at most $|L|$ (otherwise, the same location would be present twice and the remaining play would still contain a cycle). This ensures that the total cost of all transitions is bounded by $|L|\Pi^{\text{tr}} - \lfloor (|\tilde{\rho}| - 1)/|L| \rfloor$. Hence the bound given above.

Then, we consider the general case of a finite play $\rho \in \text{Play}(\sigma_{\text{Min}})$ that might cross several intervals. We achieve this by splitting the play along intervals of $\text{int}(\sigma_{\text{Min}})$. Let I_1, I_2, \dots, I_k be the intervals of $\text{int}(\sigma_{\text{Min}})$ visited during ρ (with $k \leq |\sigma_{\text{Min}}|$). We split ρ into k plays $\rho_1, \rho_2, \dots, \rho_k$ such that $\rho = \rho_1 \xrightarrow{c_1} \rho_2 \xrightarrow{c_2} \dots \rho_k$; and, for all i , all clock values along ρ_i are in I_i (remember that SPTGs contain no reset transitions). Then, we have:

$$\text{Cost}(\rho) = \sum_{i=1}^k \text{Cost}(\rho_i) + \sum_{i=1}^{k-1} c_i. \quad (3.4)$$

Let us bound these two sums. We start with the rightmost one. Since $c_i \leq \Pi^{\text{tr}}$ for all i , and since $k \leq |\sigma_{\text{Min}}|$, we have:

$$\sum_{i=1}^{k-1} c_i \leq (k-1)\Pi^{\text{tr}} \leq (|\sigma_{\text{Min}}| - 1)\Pi^{\text{tr}}. \quad (3.5)$$

Now let us bound the leftmost sum in (3.4). Using (3.3), we obtain:

$$\sum_{i=1}^k \text{Cost}(\rho_i) \leq \Pi^{\text{loc}} \sum_{i=1}^k |I_i| + \sum_{i=1}^k |L|\Pi^{\text{tr}} - \sum_{i=1}^k \lfloor (|\rho_i| - 1)/|L| \rfloor \quad (3.6)$$

Now, we can further bound these three new sums. The intervals I_i are consecutive, hence $\sum_{i=1}^k |I_i| \leq 1$. Next, $\sum_{i=1}^k |L|\Pi^{\text{tr}} = k|L|\Pi^{\text{tr}}$. But since $k \leq |\sigma_{\text{Min}}|$, we obtain

that $\sum_{i=1}^k |L|\Pi^{\text{tr}} \leq |\sigma_{\text{Min}}||L|\Pi^{\text{tr}}$. For the last sum, we observe that, by definition of the split of ρ into $\rho_1, \rho_2, \dots, \rho_k$ (with $k-1$ extra transitions in-between), $|\rho| = \sum_{i=1}^k |\rho_i|$, hence $\sum_{i=1}^k \lfloor (|\rho_i| - 1)/|L| \rfloor \geq \frac{|\rho|}{|L|} - 2|\sigma_{\text{Min}}|$, since $|\sigma_{\text{Min}}| \geq k$. Plugging these three bounds in (3.6), we obtain:

$$\sum_{i=1}^k \text{Cost}(\rho_i) \leq \Pi^{\text{loc}} + |\sigma_{\text{Min}}||L|\Pi^{\text{tr}} - \frac{|\rho|}{|L|} + 2|\sigma_{\text{Min}}|. \quad (3.7)$$

Finally, using the bounds (3.5) and (3.7) in (3.4), we obtain:

$$\text{Cost}(\rho) \leq \Pi^{\text{loc}} + (|\sigma_{\text{Min}}| - 1)\Pi^{\text{tr}} + |\sigma_{\text{Min}}||L|\Pi^{\text{tr}} - \frac{|\rho|}{|L|} + 2|\sigma_{\text{Min}}|,$$

which concludes the proof, using the fact that $|L| \geq 1$. \square

3.3.5. Fake-optimal strategies. Next, to characterise the fact that σ_{Min} must allow Min to reach a cost which is *small enough, without necessarily reaching a target location*, we define the *fake value* of an NC-strategy σ_{Min} from a configuration s as:

$$\text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}}(s) = \sup\{\text{Price}(\rho) \mid \rho \in \text{Play}(s, \sigma_{\text{Min}}), \rho \text{ reaches a target}\}$$

i.e. the value obtained when *ignoring* the σ_{Min} -induced plays that *do not* reach the target: we let $\sup \emptyset = -\infty$. Thus, clearly, $\text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}}(s) \leq \text{Val}^{\sigma_{\text{Min}}}(s)$. We say that an NC-strategy σ_{Min} is *fake-optimal* if $\text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}}(s) = \text{Val}_{\mathcal{G}}(s)$ for all configurations s .

Let us now explain why this notion of fake-optimal strategy is important. As we are about to show, we can combine any fake-optimal NC-strategy σ_{Min}^1 with an attractor strategy σ_{Min}^2 into a switching strategy σ_{Min} , which forces to eventually reach the target with a price that we can make as small as desired (since σ_{Min}^1 is an NC-strategy) when (negative) cycles can be enforced in the game by Min .

Lemma 3.9. *Let \mathcal{G} be an SPTG such that $\text{Val}_{\mathcal{G}}(s) \neq +\infty$, for all s . Let σ_{Min}^1 be an NC-strategy of Min in \mathcal{G} , and σ_{Min}^2 be an attractor strategy. Then, for all $n \in \mathbb{N}$, the switching strategy σ_{Min} described by the pair $(\sigma_{\text{Min}}^1, \sigma_{\text{Min}}^2)$ and the switching threshold*

$$K = |L| \times (2\Pi^{\text{loc}} + (2|\sigma_{\text{Min}}^1|) \times |L|\Pi^{\text{tr}} + 3|\sigma_{\text{Min}}^1| - \max(-n, \text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}^1}(s)))$$

is such that $\text{Val}_{\mathcal{G}}^{\sigma_{\text{Min}}}(s) \leq \max(-n, \text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}^1}(s))$ for all configurations s .

Remark 3.10. In particular, in the case where σ_{Min}^1 is a fake-optimal NC-strategy, and $\text{Val}_{\mathcal{G}}(s) \neq -\infty$, and when we choose the parameter n in the definition of the threshold s.t. $n > -\text{Val}_{\mathcal{G}}(s)$, then, we obtain a strategy σ_{Min}^1 that is optimal for Min from configuration s .

Proof of Lemma 3.9. In order to establish that $\text{Val}_{\mathcal{G}}^{\sigma_{\text{Min}}}(s) \leq \max(-n, \text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}^1}(s))$ (under the assumptions of the lemma), we will consider any play ρ in $\text{Play}(s, \sigma_{\text{Min}})$ and show that $\text{Price}(\rho) \leq \max(-n, \text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}^1}(s))$.

There are two possibilities regarding $\rho \in \text{Play}(s, \sigma_{\text{Min}})$, depending on whether the switch has happened or not:

- (1) If ρ reaches the target without switching from σ_{Min}^1 to σ_{Min}^2 , then $\rho \in \text{Play}(s, \sigma_{\text{Min}}^1)$ and thus $\text{Price}(\rho) \leq \text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}^1}(s) \leq \max(-n, \text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}^1}(s))$.

- (2) If ρ reaches the target after the switch happened from σ_{Min}^1 to σ_{Min}^2 , we can decompose ρ into the concatenation of a prefix ρ_1 of length $K + 1$ conforming to σ_{Min}^1 , and a play prefix ρ_2 conforming to σ_{Min}^2 . As σ_{Min}^1 is an NC-strategy, every $|L|$ steps, either the time valuation of the play changed of interval of $\text{int}(\sigma_{\text{Min}}^1)$ or a cycle occurred within the same interval in which the cost of the discrete transitions is at most -1 . In other words, if $(m + |\sigma_{\text{Min}}^1|)|L|$ steps occurred, at least m cycle occurred reducing the discrete cost of the play by at least m . As a consequence, thanks to Lemma 3.8, since ρ_1 has K steps with

$$K = |L| \times (\Pi^{\text{loc}} + (2|\sigma_{\text{Min}}^1| - 1) \times |L|\Pi^{\text{tr}} + 3|\sigma_{\text{Min}}^1| - [\max(-n, \text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}^1}(s)) - \Pi^{\text{loc}} - |L|\Pi^{\text{tr}}])$$

we know that $\text{Cost}(\rho_1) \leq \max(-n, \text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}^1}(s)) - \Pi^{\text{loc}} - |L|\Pi^{\text{tr}}$. Moreover, $\text{Cost}(\rho_2) \leq \Pi^{\text{loc}} + |L|\Pi^{\text{tr}}$ since σ_{Min}^2 follows an attractor computation and must thus reach the target in at most $|L|$ transitions (and at most 1 unit of time). Hence,

$$\text{Price}(\rho) = \text{Cost}\rho_1 + \text{Cost}\rho_2 \leq \max(-n, \text{fake}_{\mathcal{G}}^{\sigma_{\text{Min}}^1}(s)). \quad \square$$

This result allows us to identify the conditions we need to check to make sure than an SPTG admits optimal strategies that can be described in a *finite way*. Formally, we say that an SPTG is *finitely optimal* if:

- (1) Min has a fake-optimal NC-strategy;
- (2) Max has an optimal FP-strategy; and
- (3) $\text{Val}_{\mathcal{G}}(\ell)$ is a cost function, for all locations ℓ .

The central point in establishing Theorem 3.4 will thus be to prove that **all SPTGs are finitely optimal** (Theorem 5.9), as this guarantees the existence of well-behaved optimal strategies and value functions. We will also show that these have a pseudo-polynomial number of cutpoints (Theorem 5.13), which easily induces that they can be computed in pseudo-polynomial time. The proof is by induction on the number of non-urgent locations of the SPTG. In Section 4, we address the base case of SPTGs with urgent locations only (where no time can elapse). Since these SPTGs are very close to the *untimed* min-cost reachability games of [BGHM16], we adapt the algorithm in this work and obtain the `solveInstant` function (Algorithm 1). This function can also compute $\text{Val}_{\mathcal{G}}(\ell, 1)$ for all ℓ and all games \mathcal{G} (even with non-urgent locations) since time cannot elapse anymore when the clock has value 1. Next, using the continuity result of Theorem 3.2, we can detect locations ℓ where $\text{Val}_{\mathcal{G}}(\ell, \nu) \in \{+\infty, -\infty\}$, for all $\nu \in [0, 1]$, and remove them from the game. Finally, in Section 5 we handle SPTGs with non-urgent locations by refining the technique of [BLMR06, Rut11] (that work only on SPTGs with non-negative weights).

4. SPTGS WITH ONLY URGENT LOCATIONS

Throughout this section, we consider an r -SPTG $\mathcal{G} = (L_{\text{Min}}, L_{\text{Max}}, L_f, L_u, \varphi, \Delta, \pi)$ where all non-final locations are urgent, i.e. $L_u = L_{\text{Min}} \cup L_{\text{Max}}$. We also fix an initial clock value ν . Since all locations in \mathcal{G} are urgent, no time will elapse, all configurations will have the same clock value ν and no cost will be incurred by staying in the different locations of the plays. Hence, we can simplify their representation: in this section, a play $\rho = (\ell_0, \nu) \xrightarrow{c_0} (\ell_1, \nu) \xrightarrow{c_1} \dots$ will be represented simply as $\ell_0 \ell_1 \dots$. The price of this play is $\text{Price}(\rho) = +\infty$ if $\ell_k \notin L_f$ for all $k \geq 0$; and $\text{Price}(\rho) = \sum_{i=0}^{k-1} \pi(\ell_i, \ell_{i+1}) + \varphi_{\ell_k}(\nu)$ if k is the least position such that $\ell_k \in L_f$.

4.1. Computing the game value for a particular clock value. We first explain how we can compute the value function of the game for a *fixed* clock value $\nu \in [0, r]$: more precisely, we will compute the vector $(\text{Val}(\ell, \nu))_{\ell \in L}$ of values for all locations. We will denote by $\text{Val}_\nu(\ell)$ the value $\text{Val}(\ell, \nu)$, so that Val_ν is the vector we want to compute. Since no time can elapse, it consists in an adaptation of the techniques developed in [BGHM16] to solve (untimed) *min-cost reachability games*. The main difference concerns the weights being rational (and not integers) and the presence of final cost functions.

Following the arguments of [BGHM16], we first observe that locations ℓ with values $\text{Val}_\nu(\ell) = +\infty$ and $\text{Val}_\nu(\ell) = -\infty$ can be pre-computed (using respectively attractor and mean-payoff techniques) and removed from the game without changing the other values. Then, because of the particular structure of the game \mathcal{G} (where a real cost is paid only on the target location, all other weights being integers), for all plays ρ , $\text{Price}(\rho)$ is a value from the set $\mathbb{Z}_{\nu, \varphi} = \mathbb{Z} + \{\varphi_\ell(\nu) \mid \ell \in L_f\}$. We further define $\mathbb{Z}_{\nu, \varphi}^{+\infty} = \mathbb{Z}_{\nu, \varphi} \cup \{+\infty\}$. Clearly, $\mathbb{Z}_{\nu, \varphi}$ contains at most $|L_f|$ values between two consecutive integers, i.e.

$$\forall i \in \mathbb{Z} \quad |[i, i+1) \cap \mathbb{Z}_{\nu, \varphi}| \leq |L_f| \quad (4.1)$$

Then, we define an operator $\mathcal{F}: (\mathbb{Z}_{\nu, \varphi}^{+\infty})^L \rightarrow (\mathbb{Z}_{\nu, \varphi}^{+\infty})^L$ mapping every vector $\mathbf{x} = (x_\ell)_{\ell \in L}$ of $(\mathbb{Z}_{\nu, \varphi}^{+\infty})^L$ to $\mathcal{F}(\mathbf{x}) = (\mathcal{F}(\mathbf{x})_\ell)_{\ell \in L}$ defined by

$$\mathcal{F}(\mathbf{x})_\ell = \begin{cases} \varphi_\ell(\nu) & \text{if } \ell \in L_f \\ \max_{(\ell, \ell') \in \Delta} (\pi(\ell, \ell') + \mathbf{x}_{\ell'}) & \text{if } \ell \in L_{\text{Max}} \\ \min_{(\ell, \ell') \in \Delta} (\pi(\ell, \ell') + \mathbf{x}_{\ell'}) & \text{if } \ell \in L_{\text{Min}}. \end{cases}$$

We will obtain Val_ν as the limit of the sequence $(\mathbf{x}^{(i)})_{i \geq 0}$ defined by $\mathbf{x}_\ell^{(0)} = +\infty$ if $\ell \notin L_f$, and $\mathbf{x}_\ell^{(0)} = \varphi_\ell(\nu)$ if $\ell \in L_f$, and then $\mathbf{x}^{(i)} = \mathcal{F}(\mathbf{x}^{(i-1)})$ for $i \geq 1$.

The intuition behind this sequence is that $\mathbf{x}^{(i)}$ is the value of the game (when the clock takes value ν) if we impose that Min must reach the target within i steps (and pays a price of $+\infty$ if it fails to do so). Formally, for a play $\rho = \ell_0 \ell_1 \dots$, we let $\text{Price}^{\leq i}(\rho) = \text{Price}(\rho)$ if $\ell_k \in L_f$ for some $k \leq i$, and $\text{Price}^{\leq i}(\rho) = +\infty$ otherwise. We further let

$$\overline{\text{Val}}_\nu^{\leq i}(\ell) = \inf_{\sigma_{\text{Min}}} \sup_{\sigma_{\text{Max}}} \text{Price}^{\leq i}(\text{Play}((\ell, \nu), \sigma_{\text{Max}}, \sigma_{\text{Min}}))$$

where σ_{Min} and σ_{Max} are respectively strategies of Min and Max. Lemma 6 of [BGHM16] allows us to easily obtain that:

Lemma 4.1. *For all $i \geq 0$, and $\ell \in L$: $\mathbf{x}_\ell^{(i)} = \overline{\text{Val}}_\nu^{\leq i}(\ell)$.*

Sketch of proof. This is proved by induction on i . It is trivial for $i = 0$, and playing one more step amounts to computing one more iterate of \mathcal{F} . \square

Now, let us study how the sequence $(\overline{\text{Val}}_\nu^{\leq i})_{i \geq 0}$ behaves and converges to the finite values of the game. Using again the same arguments as in [BGHM16] (in particular, that \mathcal{F} is a monotonic and Scott-continuous operator over the complete lattice $(\mathbb{Z}_{\nu, \varphi}^{+\infty})^L$), the sequence $(\overline{\text{Val}}_\nu^{\leq i})_{i \geq 0}$ converges towards the greatest fixed point of \mathcal{F} . Let us now show that Val_ν is actually this greatest fixed point. First, Lemma 7 of [BGHM16] can be adapted to obtain

Lemma 4.2. *For all $\ell \in L$: $\overline{\text{Val}}_\nu^{\leq |L|-1}(\ell) \leq (|L| - 1)\Pi^{\text{tr}} + \Pi^{\text{fin}}$.*

Algorithm 1: solveInstant(\mathcal{G}, ν)

Input: r -SPTG $\mathcal{G} = (L_{\text{Min}}, L_{\text{Max}}, L_f, L_u, \varphi, \Delta, \pi)$, a clock value $\nu \in [0, r]$

```

1 foreach  $\ell \in L$  do
2   if  $\ell \in L_f$  then  $X(\ell) := \varphi_\ell(\nu)$  else  $X(\ell) := +\infty$ 
3 repeat
4    $X_{pre} := X$ 
5   foreach  $\ell \in L_{\text{Max}}$  do  $X(\ell) := \max_{(\ell, \ell') \in \Delta} (\pi(\ell, \ell') + X_{pre}(\ell'))$ 
6   foreach  $\ell \in L_{\text{Min}}$  do  $X(\ell) := \min_{(\ell, \ell') \in \Delta} (\pi(\ell, \ell') + X_{pre}(\ell'))$ 
7   foreach  $\ell \in L$  such that  $X(\ell) < -(|L| - 1)\Pi^{\text{tr}} - \Pi^{\text{fin}}$  do  $X(\ell) := -\infty$ 
8 until  $X = X_{pre}$ 
9 return  $X$ 

```

Proof. Denoting by $\text{Attr}_i(S)$ the i -steps attractor of set S (i.e. the set of locations where player Min can enforce reaching S in at most i steps), and assuming that $\text{Attr}_{-1}(S) = \emptyset$ for all S , we can establish by induction on j that: for all locations $\ell \in L$ with $0 \leq k \leq |L|$ such that $\ell \in \text{Attr}_k(L_f) \setminus \text{Attr}_{k-1}(L_f)$, and for all $0 \leq j \leq |L|$:

- (1) $j < k$ implies $\overline{\text{Val}}_\nu^{\leq j}(\ell) = +\infty$ and
- (2) $j \geq k$ implies $\overline{\text{Val}}_\nu^{\leq j}(\ell) \leq j\Pi^{\text{tr}} + \Pi^{\text{fin}}$ and $\overline{\text{Val}}_\nu^{\leq j}(\ell) \in \mathbb{Z}_{\nu, \varphi}$.

Then, the result is obtained by taking $j = |L| - 1$ in 2. □

The next step is to show that the values that can be computed along the sequence (still assuming that $\text{Val}(\ell, \nu)$ is finite for all ℓ) are taken from a finite set:

Lemma 4.3. *For all $i \geq 0$ and for all $\ell \in L$:*

$$\overline{\text{Val}}_\nu^{\leq |L|+i}(\ell) \in \text{PossVal}_\nu = [-(|L| - 1)\Pi^{\text{tr}} - \Pi^{\text{fin}}, (|L| - 1)\Pi^{\text{tr}} + \Pi^{\text{fin}}] \cap \mathbb{Z}_{\nu, \varphi}$$

where PossVal_ν has cardinality bounded by $|L_f| \times (2(|L| - 1)\Pi^{\text{tr}} + 2\Pi^{\text{fin}} + 1)$.

Proof. Following the proof of [BGHM16, Lemma 9], it is easy to show that if Min can secure, from some vertex ℓ , a price less than $-(|L| - 1)\Pi^{\text{tr}} - \Pi^{\text{fin}}$, i.e. $\text{Val}(\ell, \nu) < -(|L| - 1)\Pi^{\text{tr}} - \Pi^{\text{fin}}$, then it can secure an arbitrarily small price from that configuration, i.e. $\text{Val}(\ell, \nu) = -\infty$, which contradicts our hypothesis that the value is finite.

Hence, for all $i \geq 0$, for all ℓ : $\overline{\text{Val}}_\nu^{\leq i}(\ell) \geq \text{Val}(\ell, \nu) > -(|L| - 1)\Pi^{\text{tr}} - \Pi^{\text{fin}}$. By Lemma 4.2 and since the sequence is non-increasing, we conclude that, for all $i \geq 0$ and for all $\ell \in L$:

$$-(|L| - 1)\Pi^{\text{tr}} - \Pi^{\text{fin}} < \overline{\text{Val}}_\nu^{\leq |L|+i}(\ell) \leq (|L| - 1)\Pi^{\text{tr}} + \Pi^{\text{fin}}.$$

Since all $\overline{\text{Val}}_\nu^{\leq |L|+i}(\ell)$ are also in $\mathbb{Z}_{\nu, \varphi}$, we conclude that $\overline{\text{Val}}_\nu^{\leq |L|+i}(\ell) \in \text{PossVal}_\nu$ for all $i \geq 0$. The upper bound on the size of PossVal_ν is established by equation (4.1). □

This allows us to bound the number of iterations needed for the sequence to stabilise. Indeed, at each step after the first $|L|$ steps, the value of at least one location must decrease, while remaining in a set of values that contains $2(|L| - 1)\Pi^{\text{tr}} + 2\Pi^{\text{fin}} + 1$ elements.

Corollary 4.4. *The sequence $(\overline{\text{Val}}_\nu^{\leq i})_{i \geq 0}$ stabilises after a number of steps at most $|L_f| \times |L| \times (2(|L| - 1)\Pi^{\text{tr}} + 2\Pi^{\text{fin}} + 1) + |L|$.*

Next, the proofs of [BGHM16, Lemma 10 and Corollary 11] allow us to conclude that this sequence converges towards the value Val_ν of the game (when all values are finite), which proves that the value iteration scheme of Algorithm 1 computes exactly Val_ν for all $\nu \in [0, r]$. Indeed, this algorithm also works when some values are not finite. As a corollary, we obtain a characterisation of the possible values of \mathcal{G} :

Corollary 4.5. *For all r -SPTGs \mathcal{G} with only urgent locations, locations $\ell \in L$ and values $\nu \in [0, r]$, $\text{Val}(\ell, \nu)$ is contained in the set $\text{PossVal}_\nu \cup \{-\infty, +\infty\}$ of cardinal polynomial in $|L|$, Π^{tr} , and Π^{fin} , i.e. pseudo-polynomial with respect to the size of \mathcal{G} .*

Finally, Section 3.4 of [BGHM16] explains how to compute simultaneously optimal strategies for both players. In our context, this allows us to obtain for every clock value $\nu \in [0, r]$ and location ℓ of an r -SPTG, such that $\text{Val}(\ell, \nu) \notin \{-\infty, +\infty\}$, an optimal FP-strategy for Max, and an optimal switching strategy for Min. In case of a configuration of value $-\infty$, the switching strategies built in Lemma 3.9, for all parameters $n > 0$, give a sequence of strategies of Min that ensure a value as low as possible.

4.2. Study of the complete value functions: \mathcal{G} is finitely optimal. So far, we have been able to compute $\text{Val}_{\mathcal{G}}(\ell, \nu)$ for a fixed value ν . In practice, this can be achieved by calling `solveInstant` (Algorithm 1). Clearly, running this algorithm for all possible valuations ν is not feasible, so let us now explain how we can reduce the computation of $\text{Val}_{\mathcal{G}}(\ell): \nu \in [0, r] \mapsto \text{Val}(\ell, \nu)$ (for all ℓ) to a *finite number of calls to solveInstant*. We first study a precise characterisation of these functions, in particular showing that these are cost functions of $\text{CF}_{\{[0, r]\}}$.

We first define the set $\text{F}_{\mathcal{G}}$ of affine functions over $[0, r]$ as follows:

$$\text{F}_{\mathcal{G}} = \{k + \varphi_\ell \mid \ell \in L_f \wedge k \in [-(|L| - 1)\Pi^{\text{tr}}, (|L| - 1)\Pi^{\text{tr}}] \cap \mathbb{Z}\}$$

Observe that this set is finite and that its cardinality is bounded above by $2|L|^2\Pi^{\text{tr}}$, pseudo-polynomial in the size of \mathcal{G} . Moreover, as a direct consequence of Corollary 4.5, this set contains enough information to compute the value of the game in each possible value of the clock, in the following sense:

Lemma 4.6. *For all $\ell \in L$, for all $\nu \in [0, r]$: if $\text{Val}(\ell, \nu)$ is finite, then there is $f \in \text{F}_{\mathcal{G}}$ such that $\text{Val}(\ell, \nu) = f(\nu)$.*

Using the continuity of $\text{Val}_{\mathcal{G}}$ (Theorem 3.2), this shows that all the cutpoints of $\text{Val}_{\mathcal{G}}$ are intersections of functions from $\text{F}_{\mathcal{G}}$, i.e. belong to the set of *possible cutpoints*

$$\text{PossCP}_{\mathcal{G}} = \{\nu \in [0, r] \mid \exists f_1, f_2 \in \text{F}_{\mathcal{G}} : f_1 \neq f_2 \wedge f_1(\nu) = f_2(\nu)\}.$$

This set is depicted in Figure 5 on an example. Observe that $\text{PossCP}_{\mathcal{G}}$ contains at most $|\text{F}_{\mathcal{G}}|^2 = 4|L|^4(\Pi^{\text{tr}})^2$ points (also pseudo-polynomial in the size of \mathcal{G}) since all functions in $\text{F}_{\mathcal{G}}$ are affine, and can thus intersect at most once with every other function. Moreover, $\text{PossCP}_{\mathcal{G}} \subseteq \mathbb{Q}$, since all functions of $\text{F}_{\mathcal{G}}$ take rational values in 0 and $r \in \mathbb{Q}$. Thus, for all ℓ , $\text{Val}_{\mathcal{G}}(\ell)$ is a cost function (with cutpoints in $\text{PossCP}_{\mathcal{G}}$ and pieces from $\text{F}_{\mathcal{G}}$). Since $\text{Val}_{\mathcal{G}}(\ell)$ is a piecewise affine function, we can characterise it completely by computing only its value on its cutpoints. Hence, we can reconstruct $\text{Val}_{\mathcal{G}}(\ell)$ by calling `solveInstant` on each rational clock value $\nu \in \text{PossCP}_{\mathcal{G}}$. From the optimal strategies computed along `solveInstant`, we can also reconstruct a fake-optimal NC-strategy for Min and an optimal FP-strategy for Max, hence:




Figure 5: Network of affine functions defined by $F_{\mathcal{G}}$: functions in bold are final affine functions of \mathcal{G} , whereas non-bold ones are their translations with weights $k \in [-(|L| - 1)\Pi^{\text{tr}}, (|L| - 1)\Pi^{\text{tr}}] \cap \mathbb{Z} = \{-1, 0, 1\}$. $\text{PossCP}_{\mathcal{G}}$ is the set of abscissæ of intersection points, represented by black disks.

Proposition 4.7. *Every r -SPTG \mathcal{G} with only urgent locations is finitely optimal. Moreover, for all locations ℓ , the piecewise affine function $\text{Val}_{\mathcal{G}}(\ell)$ has cutpoints in $\text{PossCP}_{\mathcal{G}}$ of cardinality $4|L|^4(\Pi^{\text{tr}})^2$, pseudo-polynomial in the size of \mathcal{G} .*

Let us establish this proposition. Notice, that it allows us to compute $\text{Val}(\ell)$ for every $\ell \in L$. First, we compute the set $\text{PossCP}_{\mathcal{G}} = \{y_1, y_2, \dots, y_k\}$, which can be done in pseudo-polynomial time in the size of \mathcal{G} . Then, for all $1 \leq i \leq k$, we can compute the vectors $(\text{Val}(\ell, y_i))_{\ell \in L}$ of values in each location when the clock takes value y_i using Algorithm 1. This provides the value of $\text{Val}(\ell)$ in each cutpoint, for all locations ℓ , which is sufficient to characterise the whole value function, as it is continuous and piecewise affine. Observe that all cutpoints, and values at the cutpoints, in the value function are rational numbers, so Algorithm 1 is effective. Thanks to the above discussions, this procedure consists in a pseudo-polynomial number of calls to a pseudo-polynomial algorithm, hence, it runs in pseudo-polynomial time. This allows us to conclude that $\text{Val}_{\mathcal{G}}(\ell)$ is a cost function for all ℓ . This proves item 3 of the definition of finite optimality for r -SPTGs with only urgent locations.

Let us conclude the proof that r -SPTGs with only urgent locations are finitely optimal by showing that Min has a fake-optimal NC-strategy, and Max has an optimal FP-strategy. Let $\nu_1, \nu_2, \dots, \nu_k$ be the sequence of elements from $\text{PossCP}_{\mathcal{G}}$ in increasing order, and let us assume $\nu_0 = 0$. For all $1 \leq i \leq k$, let f_i^ℓ be the function from $F_{\mathcal{G}}$ that defines the piece of $\text{Val}_{\mathcal{G}}(\ell)$ in the interval $[\nu_{i-1}, \nu_i]$ (we have shown above that such an f_i^ℓ always exists). Formally, for all $1 \leq i \leq k$, $f_i^\ell \in F_{\mathcal{G}}$ verifies $\text{Val}(\ell, \nu) = f_i^\ell(\nu)$, for all $\nu \in [\nu_{i-1}^\ell, \nu_i^\ell]$. Next, for all $1 \leq i \leq k$, let μ_i be a value taken in the middle⁶ of $[\nu_{i-1}, \nu_i]$, i.e. $\mu_i = \frac{\nu_i + \nu_{i-1}}{2}$. Note that all μ_i 's are rational values since all ν_i 's are. By applying `solveInstant` in each μ_i , we can compute $(\text{Val}_{\mathcal{G}}(\ell, \mu_i))_{\ell \in L}$, and we can extract an optimal memoryless strategy σ_{Max}^i

⁶Taking the middle is an arbitrary choice, any point strictly within the interval would work

for Max and an optimal switching strategy σ_{Min}^i for Min. Thus we know that, for all $\ell \in L$, playing σ_{Min}^i (respectively, σ_{Max}^i) from (ℓ, μ_i) allows Min (respectively, Max) to ensure a price at most (respectively, at least) $\text{Val}_{\mathcal{G}}(\ell, \mu_i) = f_i^\ell(\mu_i)$. However, it is easy to check that the bound given by $f_i^\ell(\mu_i)$ holds in every clock value, i.e. for all ℓ , for all $\nu \in [\nu_{i-1}, \nu_i]$

$$\text{Price}((\ell, \nu), \sigma_{\text{Min}}^i) \leq f_i^\ell(\nu) \quad \text{and} \quad \text{Price}((\ell, \nu), \sigma_{\text{Max}}^i) \geq f_i^\ell(\nu).$$

This holds because:

- (1) Min can play σ_{Min}^i from all clock values (in $[0, r]$) since we are considering an r -SPTG; and
- (2) Max does not have more possible strategies from an arbitrary clock value $\nu \in [0, r]$ than from μ_i , because all locations are urgent and time cannot elapse (neither from ν , nor from μ_i).

And symmetrically for Max.

We conclude that Min can consistently play the same strategy σ_{Min}^i from all configurations (ℓ, ν) with $\nu \in [\nu_{i-1}, \nu_i]$ and secure a price which is at most $f_i^\ell(\nu) = \text{Val}_{\mathcal{G}}(\ell, \nu)$, i.e. σ_{Min}^i is optimal on this interval. By definition of σ_{Min}^i , it is easy to extract from it a fake-optimal NC-strategy (actually, σ_{Min}^i is a switching strategy described by a pair $(\sigma_{\text{Min}}^1, \sigma_{\text{Min}}^2)$, and σ_{Min}^1 can be used to obtain the fake-optimal NC-strategy). The same reasoning applies to strategies of Max and we conclude that Max has an optimal FP-strategy.

5. FINITE OPTIMALITY OF GENERAL SPTGS

In this section, we consider SPTGs with non-urgent locations. We first prove that all such SPTGs are finitely optimal. Then, we introduce Algorithm 2 to compute optimal values and strategies of SPTGs. Throughout the section, we fix an SPTG $\mathcal{G} = (L_{\text{Min}}, L_{\text{Max}}, L_f, L_u, \varphi, \Delta, \pi)$ with non-urgent locations. Before presenting our core contributions, let us explain how we can detect locations with infinite values. As already argued, we can compute $\text{Val}(\ell, 1)$ for all ℓ assuming all locations are urgent, since time cannot elapse anymore when the clock has value 1. This can be done with `solveInstant` (Algorithm 1). Then, from the absence of guards in SPTGs and Theorem 3.2 we have that $\text{Val}(\ell, 1) = +\infty$ (respectively, $\text{Val}(\ell, 1) = -\infty$) if and only if $\text{Val}(\ell, \nu) = +\infty$ (respectively, $\text{Val}(\ell, \nu) = -\infty$) for all $\nu \in [0, 1]$. We can thus remove from the game all locations with infinite value, and this does not affect the values of other locations. Thus, we henceforth assume that $\text{Val}(\ell, \nu) \in \mathbb{R}$ for all $(\ell, \nu) \in \text{Conf}_{\mathcal{G}}$.

5.1. The $\mathcal{G}_{L',r}$ construction. To prove finite optimality of SPTGs and to establish correctness of our algorithm, we rely in both cases on a construction that consists in decomposing \mathcal{G} into a sequence of SPTGs with *fewer non-urgent locations*. Intuitively, a game with fewer non-urgent locations is easier to solve since it is closer to an untimed game (in particular, when all locations are urgent, we can apply the techniques of Section 4). More precisely, given a set L' of non-urgent locations, we will define a (possibly infinite) sequence of clock values $1 = r_0 > r_1 > \dots$ and a sequence $\mathcal{G}_{L',r_0}, \mathcal{G}_{L',r_1}, \dots$ of SPTGs such that

- (1) the non-final locations of \mathcal{G}_{L',r_i} are exactly the ones of \mathcal{G} , except that the locations of L' are now urgent (some final locations are added to allow players to wait until r_i); and

(2) for all $i \geq 0$, the value function of \mathcal{G}_{L',r_i} is equal to $\text{Val}_{\mathcal{G}}$ on the interval $[r_{i+1}, r_i]$. Hence, we can re-construct $\text{Val}_{\mathcal{G}}$ by assembling well-chosen parts of the value functions of the games \mathcal{G}_{L',r_i} (assuming $\inf_i r_i = 0$).

In fact, we will show later (see Lemma 5.8) that we can assume the sequence r_0, \dots to be finite. This basic result will be exploited in two directions. First, we prove by induction on the number of non-urgent locations that all SPTGs are finitely optimal, by re-constructing $\text{Val}_{\mathcal{G}}$ (as well as optimal strategies) as a \triangleright -concatenation of the value functions of a finite sequence of SPTGs with one non-urgent locations less. The base case, with only urgent locations, is solved by Proposition 4.7. This construction suggests a *recursive* algorithm in the spirit of [BLMR06, Rut11] (for non-negative weights). Second, we show that this recursion can be *avoided* (see Algorithm 2). Instead of turning locations urgent one at a time, this algorithm makes them all urgent and computes directly the sequence of SPTGs with only urgent locations. Its proof of correctness relies on the finite optimality of SPTGs and, again, on our basic result linking the value functions of \mathcal{G} and games \mathcal{G}_{L',r_i} .

Let us formalise these constructions. Let \mathcal{G} be an SPTG, $r \in [0, 1]$ be an endpoint, and $\mathbf{x} = (x_\ell)_{\ell \in L}$ be a vector of rational values. Then, $\text{wait}(\mathcal{G}, r, \mathbf{x})$ is an r -SPTG in which both players may now decide, in all non-urgent locations ℓ , to *wait* until the clock takes value r , and then to stop the game, adding the weight x_ℓ to the current cost of the play. Formally, $\text{wait}(\mathcal{G}, r, \mathbf{x}) = (L_{\text{Min}}, L_{\text{Max}}, L'_f, L_u, \varphi', T', \pi')$ is such that

- $L'_f = L_f \uplus \{\ell^f \mid \ell \in L \setminus (L_u \cup L_f)\}$;
- for all $\ell' \in L_f$ and $\nu \in [0, r]$, $\varphi'_{\ell'}(\nu) = \varphi_{\ell'}(\nu)$, for all $\ell \in L \setminus (L_u \cup L_f)$, $\varphi'_{\ell^f}(\nu) = (r - \nu) \cdot \pi(\ell) + x_\ell$;
- $T' = T \cup \{(\ell, [0, r], \perp, \ell^f) \mid \ell \in L \setminus (L_u \cup L_f)\}$;
- for all $\delta \in T'$, $\pi'(\delta) = \pi(\delta)$ if $\delta \in T$, and $\pi'(\delta) = 0$ otherwise.

Then, we let $\mathcal{G}_r = \text{wait}(\mathcal{G}, r, (\text{Val}_{\mathcal{G}}(\ell, r))_{\ell \in L})$, i.e. the game obtained thanks to *wait* by letting \mathbf{x} be the value of \mathcal{G} in r . It is easy to check that this first transformation does not alter the value of the game, for clock values before r :

Lemma 5.1. *For all $\nu \in [0, r]$ and locations ℓ , $\text{Val}_{\mathcal{G}}(\ell, \nu) = \text{Val}_{\mathcal{G}_r}(\ell, \nu)$.*

Next, we make locations urgent. For a set $L' \subseteq L \setminus (L_u \cup L_f)$ of non-urgent locations, we let $\mathcal{G}_{L',r}$ be the SPTG obtained from \mathcal{G}_r by making urgent every location ℓ of L' . Observe that, although all locations $\ell \in L'$ are now urgent in $\mathcal{G}_{L',r}$, their clones ℓ^f allow the players to wait until r . When L' is a singleton $\{\ell\}$, we write $\mathcal{G}_{\ell,r}$ instead of $\mathcal{G}_{\{\ell\},r}$.

While the construction of \mathcal{G}_r does not change the value of the game, turning locations urgent *does*. Yet, we can characterise an interval $[a, r]$ on which the value functions of $\mathcal{H} = \mathcal{G}_{L',r}$ and $\mathcal{H}^+ = \mathcal{G}_{L' \cup \{\ell\},r}$ coincide, as stated by the next proposition. The interval $[a, r]$ depends on the *slopes* of the pieces of $\text{Val}_{\mathcal{H}^+}$ as depicted in Figure 6: for each location ℓ of Min, the slopes of the pieces of $\text{Val}_{\mathcal{H}^+}$ contained in $[a, r]$ should be $\geq -\pi(\ell)$ (and $\leq -\pi(\ell)$ when ℓ belongs to Max). It is proved by lifting optimal strategies of \mathcal{H}^+ into \mathcal{H} , and strongly relies on the determinacy result of Theorem 2.2. Hereafter, we denote the slope of $\text{Val}_{\mathcal{G}}(\ell)$ in-between ν and ν' by $\text{slope}_{\mathcal{G}}^{\ell}(\nu, \nu')$, formally defined by $\text{slope}_{\mathcal{G}}^{\ell}(\nu, \nu') = \frac{\text{Val}_{\mathcal{G}}(\ell, \nu') - \text{Val}_{\mathcal{G}}(\ell, \nu)}{\nu' - \nu}$.

Proposition 5.2. *Let $0 \leq a < r \leq 1$, $L' \subseteq L \setminus (L_u \cup L_f)$ and $\ell \notin L' \cup L_u$ a non-urgent location of Min (respectively, Max). Assume that $\mathcal{G}_{L' \cup \{\ell\},r}$ is finitely optimal, and that, for all $a \leq \nu_1 < \nu_2 \leq r$:*

$$\text{slope}_{\mathcal{G}_{L' \cup \{\ell\},r}}^{\ell}(\nu_1, \nu_2) \geq -\pi(\ell) \quad (\text{respectively, } \leq -\pi(\ell)). \quad (5.1)$$




Figure 6: The condition (5.1) (in the case $L' = \emptyset$ and $\ell \in L_{\text{Min}}$): graphically, it means that the slope between every two points of the plot in $[a, r]$ (represented with a thick line) is greater than or equal to $-\pi(\ell)$ (represented with a dashed line). The value function is depicted here as a non-piecewise-affine function, as this is not the crucial property we want to highlight.

Then, for all $\nu \in [a, r]$ and $\ell' \in L$, $\text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell', \nu) = \text{Val}_{\mathcal{G}_{L', r}}(\ell', \nu)$. Furthermore, fake-optimal NC-strategies and optimal FP-strategies in $\mathcal{G}_{L' \cup \{\ell\}, r}$ are also fake-optimal and optimal over $[a, r]$ in $\mathcal{G}_{L', r}$.

Before proving this result, we start with an auxiliary lemma showing a property of the rates of change of the value functions associated to non-urgent locations

Lemma 5.3. *Let \mathcal{G} be an r -SPTG, ℓ and ℓ' be non-urgent locations of Min and Max, respectively. Then for all $0 \leq \nu < \nu' \leq r$:*

$$\text{slope}_{\mathcal{G}}^{\ell}(\nu, \nu') \geq -\pi(\ell) \quad \text{and} \quad \text{slope}_{\mathcal{G}}^{\ell'}(\nu, \nu') \leq -\pi(\ell').$$

Proof. For the location ℓ , the inequality rewrites in

$$\text{Val}_{\mathcal{G}}(\ell, \nu) \leq (\nu' - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, \nu').$$

Using the upper definition of the value (thanks to the determinacy result of Theorem 2.2), it suffices to prove, for all $\varepsilon > 0$, the existence of a strategy σ_{Min} of Min such that for all strategies σ_{Max} of Max:

$$\text{Price}(\text{CPlay}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) \leq (\nu' - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, \nu') + \varepsilon. \quad (5.2)$$

To prove the existence of such a σ_{Min} , we first fix, given ε , a strategy σ'_{Min} such that for all strategies σ_{Max} :

$$\text{Price}(\text{CPlay}((\ell, \nu'), \sigma'_{\text{Min}}, \sigma_{\text{Max}})) \leq \text{Val}_{\mathcal{G}}(\ell, \nu') + \varepsilon.$$

Such a strategy necessarily exists by definition of the value. Then, σ_{Min} can be obtained as follows. Under σ_{Min} , Min will all always play as indicated by σ'_{Min} , except in the first round.

In this first round, the game is still in ℓ and Min will play like σ'_{Min} , adding an extra delay of $\nu' - \nu$ time units (observe Min is allowed to do so, since ℓ is non-urgent). Clearly, this extra delay in ℓ will incur a cost of $(\nu' - \nu)\pi(\ell)$, hence, we obtain (5.2).

A similar reasoning allows us to obtain the result for ℓ' . \square

Now, we show that, even if the locations in L' are turned into urgent locations, we may still obtain for them a similar result of the rates of change as the one of Lemma 5.3:

Lemma 5.4. *For all locations $\ell \in L' \cap L_{\text{Min}}$ (respectively, $\ell \in L' \cap L_{\text{Max}}$), and $\nu \in [0, r]$, $\text{Val}_{\mathcal{G}_{L',r}}(\ell, \nu) \leq (r - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r)$ (respectively, $\text{Val}_{\mathcal{G}_{L',r}}(\ell, \nu) \geq (r - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r)$).*

Proof. It suffices to notice that from (ℓ, ν) , Min (respectively, Max) may choose to go directly in ℓ^f ensuring the value $(r - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r)$. \square

We are now ready to establish Proposition 5.2:

Proof of Proposition 5.2. Let σ_{Min} and σ_{Max} be respectively a fake-optimal NC-strategy of Min and an optimal FP-strategy of Max in $\mathcal{G}_{L' \cup \{\ell\}, r}$. Notice that both strategies are also well-defined finite positional strategies in $\mathcal{G}_{L', r}$.

First, let us show that σ_{Min} is indeed an NC-strategy in $\mathcal{G}_{L', r}$. Take a finite play $(\ell_0, \nu_0) \xrightarrow{c_0} \dots \xrightarrow{c_{k-1}} (\ell_k, \nu_k)$, of length $k \geq 2$, that conforms with σ_{Min} in $\mathcal{G}_{L', r}$, and with $\ell_0 = \ell_k$ and ν_0, ν_k in the same interval I of $\text{int}(\sigma_{\text{Min}})$. To show that σ_{Min} is an NC-strategy, we need to show that the total cost of the transitions in this play is at most -1 . As σ_{Min} is finite positional and ν_0 and ν_k are in the same interval, the play $(\ell_0, \nu_k) \xrightarrow{c'_0} \dots \xrightarrow{c'_{k-1}} (\ell_k, \nu_k)$ also conforms with σ_{Min} (with possibly different weights). Furthermore, as all the delays in this new play are 0 we are sure that this play is also a valid play in $\mathcal{G}_{L' \cup \{\ell\}, r}$, in which σ_{Min} is an NC-strategy. Therefore, $\pi(\ell_0, \ell_1) + \dots + \pi(\ell_{k-1}, \ell_k) \leq -1$, and σ_{Min} is an NC-strategy in $\mathcal{G}_{L', r}$.

We now show the result for $\ell \in L_{\text{Min}}$. The proof for $\ell \in L_{\text{Max}}$ is a straightforward adaptation. Notice that every play in $\mathcal{G}_{L', r}$ that conforms with σ_{Min} is also a play in $\mathcal{G}_{L' \cup \{\ell\}, r}$ that conforms with σ_{Min} , as σ_{Min} is defined in $\mathcal{G}_{L' \cup \{\ell\}, r}$ and thus plays with no delay in location ℓ . Thus, for all $\nu \in [a, r]$ and $\ell' \in L$, by Lemma 3.9,

$$\text{Val}_{\mathcal{G}_{L', r}}(\ell', \nu) \leq \text{fake}_{\mathcal{G}_{L', r}}^{\sigma_{\text{Min}}}(\ell', \nu) = \text{fake}_{\mathcal{G}_{L' \cup \{\ell\}, r}}^{\sigma_{\text{Min}}}(\ell', \nu) = \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell', \nu). \quad (5.3)$$

To obtain that $\text{Val}_{\mathcal{G}_{L', r}}(\ell', \nu) = \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell', \nu)$, it remains to show the reverse inequality. To that extent, let ρ be a finite play in $\mathcal{G}_{L', r}$ that conforms with σ_{Max} , starts in a configuration (ℓ', ν) with $\nu \in [a, r]$, and ends in a final location. We show by induction on the length of ρ that $\text{Price}(\rho) \geq \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell', \nu)$. If ρ has size 1 then ℓ' is a final configuration and $\text{Price}(\rho) = \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell', \nu) = \varphi'_{\ell'}(\nu)$.

Otherwise $\rho = (\ell', \nu) \xrightarrow{c} \rho'$ where ρ' is a play that conforms with σ_{Max} , starting in a configuration (ℓ'', ν'') and ending in a final configuration. By induction hypothesis, we have $\text{Price}(\rho') \geq \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell'', \nu'')$. We now distinguish three cases, the two first being immediate:

- If $\ell' \in L_{\text{Max}}$, then $\sigma_{\text{Max}}(\ell', \nu)$ leads to the next configuration (ℓ'', ν'') , thus

$$\begin{aligned} \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell', \nu) &= \text{Price}_{\mathcal{G}_{L' \cup \{\ell\}, r}}((\ell', \nu), \sigma_{\text{Max}}) \\ &= c + \text{Price}_{\mathcal{G}_{L' \cup \{\ell\}, r}}((\ell'', \nu''), \sigma_{\text{Max}}) \\ &\leq c + \text{Price}(\rho') = \text{Price}(\rho). \end{aligned}$$

- If $\ell' \in L_{\text{Min}}$, and $\ell' \neq \ell$ or $\nu'' = \nu$, we have that $(\ell', \nu) \xrightarrow{c} (\ell'', \nu'')$ is a valid transition in $\mathcal{G}'_{L' \cup \{\ell\}, r}$. Therefore, $\text{Val}_{\mathcal{G}'_{L' \cup \{\ell\}, r}}(\ell', \nu) \leq c + \text{Val}_{\mathcal{G}'_{L' \cup \{\ell\}, r}}(\ell'', \nu'')$, hence

$$\text{Price}(\rho) = c + \text{Price}(\rho') \geq c + \text{Val}_{\mathcal{G}'_{L' \cup \{\ell\}, r}}(\ell'', \nu'') \geq \text{Val}_{\mathcal{G}'_{L' \cup \{\ell\}, r}}(\ell', \nu).$$

- Finally, if $\ell' = \ell$ and $\nu'' > \nu$, then $c = (\nu'' - \nu)\pi(\ell) + \pi(\ell, \ell'')$. As $(\ell, \nu'') \xrightarrow{\pi(\ell, \ell'')} (\ell'', \nu'')$ is a valid transition in $\mathcal{G}_{L' \cup \{\ell\}, r}$, we have $\text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell, \nu'') \leq \pi(\ell, \ell'') + \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell'', \nu'')$. Furthermore, since $\nu'' \in [a, r]$, we can use (5.1) to obtain

$$\begin{aligned} \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell, \nu) &\leq \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell, \nu'') + (\nu'' - \nu)\pi(\ell) \\ &\leq \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell'', \nu'') + \pi(\ell, \ell'') + (\nu'' - \nu)\pi(\ell). \end{aligned}$$

Therefore

$$\begin{aligned} \text{Price}(\rho) &= (\nu'' - \nu)\pi(\ell) + \pi(\ell, \ell'') + \text{Price}(\rho') \\ &\geq (\nu'' - \nu)\pi(\ell) + \pi(\ell, \ell'') + \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell'', \nu'') \geq \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell', \nu). \end{aligned}$$

This concludes the induction. As a consequence,

$$\inf_{\sigma'_{\text{Min}} \in \text{Strat}_{\text{Min}}(\mathcal{G}_{L', r})} \text{Price}_{\mathcal{G}_{L', r}}(\text{CPlay}((\ell', \nu), \sigma'_{\text{Min}}, \sigma_{\text{Max}})) \geq \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell', \nu)$$

for all locations ℓ' and $\nu \in [a, r]$, which finally proves that $\text{Val}_{\mathcal{G}_{L', r}}(\ell', \nu) \geq \text{Val}_{\mathcal{G}_{L' \cup \{\ell\}, r}}(\ell', \nu)$. Fake-optimality of σ_{Min} over $[a, r]$ in $\mathcal{G}_{L' \cup \{\ell\}, r}$ is then obtained by (5.3). \square

Given an SPTG \mathcal{G} and some *finitely optimal* $\mathcal{G}_{L', r}$, we now characterise precisely the left endpoint of the maximal interval ending in r where the value functions of \mathcal{G} and $\mathcal{G}_{L', r}$ coincide. To this end, we use the operator $\text{left}_{L'}: (0, 1] \rightarrow [0, 1]$ defined as:

$$\text{left}_{L'}^{\mathcal{G}}(r) = \inf\{r' \leq r \mid \forall \ell \in L \forall \nu \in [r', r] \text{Val}_{\mathcal{G}_{L', r}}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, \nu)\}.$$

Most of the time, we will forget about the \mathcal{G} exponent in $\text{left}_{L'}^{\mathcal{G}}(r)$, but we keep it since it will become useful in later proofs. By continuity of the value (Theorem 3.2), the infimum in the definition exists and $\text{Val}_{\mathcal{G}}(\ell, \text{left}_{L'}(r)) = \text{Val}_{\mathcal{G}_{L', r}}(\ell, \text{left}_{L'}(r))$. Moreover, $\text{Val}_{\mathcal{G}}(\ell)$ is a cost function on $[\text{left}_{L'}(r), r]$, since $\mathcal{G}_{L', r}$ is finitely optimal.

However, this definition of $\text{left}_{L'}(r)$ is semantical. Yet, building on the ideas of Proposition 5.2, we can effectively compute $\text{left}_{L'}(r)$, given $\text{Val}_{\mathcal{G}_{L', r}}$. We claim that $\text{left}_{L'}(r)$ is the *minimal clock value* such that for all locations $\ell \in L' \cap L_{\text{Min}}$ (respectively, $\ell \in L' \cap L_{\text{Max}}$), the slopes of the affine sections of the cost function $\text{Val}_{\mathcal{G}_{L', r}}(\ell)$ on $[\text{left}_{L'}(r), r]$ are at least (at most) $-\pi(\ell)$. Notice that while this condition (that we show formally in Lemma 5.5) only speaks about locations of L' that are made urgent, the semantical definition of $\text{left}_{L'}^{\mathcal{G}}(r)$ gives an equality of values for *all* locations of L . Via this condition, $\text{left}_{L'}(r)$ can be obtained (see Figure 7) by inspecting iteratively, for all ℓ of Min (respectively, Max), the slopes of $\text{Val}_{\mathcal{G}_{L', r}}(\ell)$, for $\ell \in L'$, by decreasing clock values until we find a piece with a slope greater than $-\pi(\ell)$ (respectively, smaller than $-\pi(\ell)$). This enumeration of the slopes is effective as

$\text{Val}_{\mathcal{G}_{L',r}}$ has finitely many pieces, by hypothesis. Moreover, this guarantees that $\text{left}_{L'}(r) < r$, as shown in the following lemma.

Lemma 5.5. *Let \mathcal{G} be an SPTG, $L' \subseteq L \setminus (L_u \cup L_f)$, and $r \in (0, 1]$, such that $\mathcal{G}_{L',r}$ is finitely optimal for all $L'' \subseteq L'$. Then, $\text{left}_{L'}(r)$ is the minimal clock value such that for all locations $\ell \in L' \cap L_{\text{Min}}$ (respectively, $\ell \in L' \cap L_{\text{Max}}$), the slopes of the affine sections of the cost function $\text{Val}_{\mathcal{G}_{L',r}}(\ell)$ on $[\text{left}_{L'}(r), r]$ are at least (respectively, at most) $-\pi(\ell)$. Moreover, $\text{left}_{L'}(r) < r$.*

Proof. Since $\text{Val}_{\mathcal{G}_{L',r}}(\ell) = \text{Val}_{\mathcal{G}}(\ell)$ on $[\text{left}_{L'}(r), r]$, and as ℓ is non-urgent in \mathcal{G} , Lemma 5.3 states that all the slopes of $\text{Val}_{\mathcal{G}}(\ell)$ are at least (respectively, at most) $-\pi(\ell)$ on $[\text{left}_{L'}(r), r]$.

We now show the minimality property by contradiction. Therefore, let $r' < \text{left}_{L'}(r)$ such that all cost functions $\text{Val}_{\mathcal{G}_{L',r}}(\ell)$ are affine on $[r', \text{left}_{L'}(r)]$, and assume that for all $\ell' \in L' \cap L_{\text{Min}}$ (respectively, $\ell' \in L' \cap L_{\text{Max}}$), the slopes of $\text{Val}_{\mathcal{G}_{L',r}}(\ell')$ on $[r', \text{left}_{L'}(r)]$ are at least (respectively, at most) $-\pi(\ell')$. Hence, this property holds on $[r', r]$. Then, by applying Proposition 5.2 $|L'|$ times (here, we use the finite optimality of the games $\mathcal{G}_{L',r}$ with $L'' \subseteq L'$), we have that for all $\nu \in [r', r]$ $\text{Val}_{\mathcal{G}_r}(\ell, \nu) = \text{Val}_{\mathcal{G}_{L',r}}(\ell, \nu)$. Using Lemma 5.1, we also know that for all $\nu \leq r$, and ℓ , $\text{Val}_{\mathcal{G}_r}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, \nu)$. Thus, $\text{Val}_{\mathcal{G}_{L',r}}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, \nu)$. As $r' < \text{left}_{L'}(r)$, this contradicts the definition of $\text{left}_{L'}(r)$.

We finally prove that $\text{left}_{L'}(r) < r$. This is immediate in case $\text{left}_{L'}(r) = 0$, since $r > 0$. Otherwise, from the result obtained previously, we know that there exists $r' < \text{left}_{L'}(r)$, and $\ell^* \in L'$ such that $\text{Val}_{\mathcal{G}_{L',r}}(\ell^*)$ is affine on $[r', \text{left}_{L'}(r)]$ of slope smaller (respectively, greater) than $-\pi(\ell^*)$ if $\ell^* \in L_{\text{Min}}$ (respectively, $\ell^* \in L_{\text{Max}}$), i.e.

$$\begin{cases} \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, r') > \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, \text{left}_{L'}(r)) + (\text{left}_{L'}(r) - r')\pi(\ell^*) & \text{if } \ell^* \in L_{\text{Min}} \\ \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, r') < \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, \text{left}_{L'}(r)) + (\text{left}_{L'}(r) - r')\pi(\ell^*) & \text{if } \ell^* \in L_{\text{Max}}. \end{cases}$$

From Lemma 5.4, we also know that

$$\begin{cases} \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, r') \leq \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, r) + (r - r')\pi(\ell^*) & \text{if } \ell^* \in L_{\text{Min}} \\ \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, r') \geq \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, r) + (r - r')\pi(\ell^*) & \text{if } \ell^* \in L_{\text{Max}}. \end{cases}$$

Both equations combined imply

$$\begin{cases} \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, r) > \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, \text{left}_{L'}(r)) + (\text{left}_{L'}(r) - r)\pi(\ell^*) & \text{if } \ell^* \in L_{\text{Min}} \\ \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, r) < \text{Val}_{\mathcal{G}_{L',r}}(\ell^*, \text{left}_{L'}(r)) + (\text{left}_{L'}(r) - r)\pi(\ell^*) & \text{if } \ell^* \in L_{\text{Max}} \end{cases}$$

which is not possible if $\text{left}_{L'}(r) = r$. □

Thus, one can reconstruct $\text{Val}_{\mathcal{G}}$ on $[\inf_i r_i, r_0]$ from the value functions of the (potentially infinite) sequence of games $\mathcal{G}_{L',r_0}, \mathcal{G}_{L',r_1}, \dots$ where $r_{i+1} = \text{left}_{L'}(r_i)$ for all i such that $r_i > 0$, for all possible choices of non-urgent locations L' . Another interesting fact, that we formally state and prove in the next lemma, is that just on the left of such a point r_i the slope of $\text{Val}_{\mathcal{G}}(\ell)$, for $\ell \in L'$, is $-\pi(\ell)$:

Lemma 5.6. *Let \mathcal{G} be an SPTG, $L' \subseteq L \setminus (L_u \cup L_f)$ and $r_0 \in (0, 1]$ such that \mathcal{G}_{L',r_0} is finitely optimal. Suppose that $r_1 = \text{left}_{L'}(r_0) > 0$, and let $r_2 = \text{left}_{L'}(r_1)$. Then, there exists $r' \in [r_2, r_1)$ and $\ell \in L'$ such that*

- (1) $\text{Val}_{\mathcal{G}}(\ell)$ is affine on $[r', r_1]$, of slope equal to $-\pi(\ell)$, and
- (2) $\text{Val}_{\mathcal{G}}(\ell, r_1) \neq \text{Val}_{\mathcal{G}}(\ell, r_0) + \pi(\ell)(r_0 - r_1)$.

As a consequence, $\text{Val}_{\mathcal{G}}(\ell)$ has a cutpoint in $[r_1, r_0]$.

Proof. We denote by r' the smallest clock value (smaller than r_1) such that for all locations ℓ , $\text{Val}_{\mathcal{G}}(\ell)$ is affine over $[r', r_1]$. Then, the proof goes by contradiction: using Lemma 5.5, we assume that for all $\ell \in L' \cap L_{\text{Min}}$ (respectively, $\ell \in L' \cap L_{\text{Max}}$)

- either $(\neg 1)$ the slope of $\text{Val}_{\mathcal{G}}(\ell)$ on $[r', r_1]$ is greater (respectively, smaller) than $-\pi(\ell)$,
- or $(1 \wedge \neg 2)$ for all $\nu \in [r', r_1]$, $\text{Val}_{\mathcal{G}}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, r_0) + \pi(\ell)(r_0 - \nu)$.

Let σ_{Min}^0 and σ_{Max}^0 (respectively, σ_{Min}^1 and σ_{Max}^1) be a fake-optimal NC-strategy and an optimal FP-strategy in \mathcal{G}_{L', r_0} (respectively, \mathcal{G}_{L', r_1}). Let $r'' = \max([r', r_1] \cap (\text{pts}(\sigma_{\text{Min}}^1) \cup \text{pts}(\sigma_{\text{Max}}^1)))$, so that strategies σ_{Min}^1 and σ_{Max}^1 have the *same behaviour* on all clock values of the interval (r'', r_1) , i.e. either always play urgently the same transition, or wait, in a non-urgent location, until reaching some clock value greater than or equal to r_1 and then play the same transition (recall that pts represent the set of endpoints in which an FP-strategy may change its behaviour).

Observe first that for all $\ell \in L' \cap L_{\text{Min}}$ (respectively, $\ell \in L' \cap L_{\text{Max}}$), if on the interval (r'', r_1) , σ_{Min}^1 (respectively, σ_{Max}^1) goes to ℓ^f then the slope on $[r'', r_1]$ (and thus on $[r', r_1]$) is $-\pi(\ell)$. This implies that (1) holds and as either $(\neg 1)$ or $(1 \wedge \neg 2)$ is true by assumption, for such a location ℓ , we know that $(1 \wedge \neg 2)$ holds (by letting $r' = r''$).

For other locations ℓ (notice that there necessarily exists a location that does not satisfy $(1 \wedge \neg 2)$ otherwise $\text{left}_{L'}(r_0) \leq r' < r_1$), we will construct a new pair of NC- and FP-strategies σ_{Min} and σ_{Max} in \mathcal{G}_{L', r_0} such that for all locations ℓ and clock values $\nu \in (r'', r_1)$

$$\text{fake}_{\mathcal{G}_{L', r_0}}^{\sigma_{\text{Min}}^0}(\ell, \nu) \leq \text{Val}_{\mathcal{G}}(\ell, \nu) \leq \text{Price}_{\mathcal{G}_{L', r_0}}((\ell, \nu), \sigma_{\text{Max}}). \quad (5.4)$$

As a consequence, with Lemma 3.9 (over game \mathcal{G}_{L', r_0}), one would have that $\text{Val}_{\mathcal{G}_{L', r_0}}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, \nu)$, which will raise a contradiction with the definition of r_1 as $\text{left}_{L'}(r_0) < r_0$ (by Lemma 5.5), and conclude the proof.

We only show the construction for σ_{Min} , as it is very similar for σ_{Max} . Strategy σ_{Min} is obtained by combining strategies σ_{Min}^1 over $[0, r_1]$, and σ_{Min}^0 over $[r_1, r_0]$: a special care has to be spent in case σ_{Min}^1 performs a jump to a location ℓ^f , since then, in σ_{Min} , we rather glue this move with the decision of strategy σ_{Min}^0 in (ℓ, r_1) . Formally, let (ℓ, ν) be a configuration of \mathcal{G}_{L', r_0} with $\ell \in L_{\text{Min}}$. We construct $\sigma_{\text{Min}}(\ell, \nu)$ as follows:

- if $\nu \geq r_1$, $\sigma_{\text{Min}}(\ell, \nu) = \sigma_{\text{Min}}^0(\ell, \nu)$;
- if $\nu < r_1$, $\ell \notin L'$ and $\sigma_{\text{Min}}^1(\ell, \nu) = (t, (\ell, \ell^f))$ for some delay t (such that $\nu + t \geq r_1$), we let $\sigma_{\text{Min}}(\ell, \nu) = (r_1 - \nu + t', (\ell, \ell'))$ where $(t', (\ell, \ell')) = \sigma_{\text{Min}}^0(\ell, r_1)$;
- otherwise $\sigma_{\text{Min}}(\ell, \nu) = \sigma_{\text{Min}}^1(\ell, \nu)$.

For all finite plays ρ in \mathcal{G}_{L', r_0} that conform to σ_{Min} , start in a configuration (ℓ, ν) such that $\nu \in (r'', r_0]$ and $\ell \notin \{\ell^f \mid \ell' \in L\}$, and end in a final location, we show by induction that $\text{Price}_{\mathcal{G}_{L', r_0}}(\rho) \leq \text{Val}_{\mathcal{G}}(\ell, \nu)$. Note that ρ either only contains clock values in $[r_1, r_0]$, or is of the form $(\ell, \nu) \xrightarrow{c} (\ell^f, \nu')$, or is of the form $(\ell, \nu) \xrightarrow{c} \rho'$ with ρ' a play that satisfies the above restriction.

- If $\nu \in [r_1, r_0]$, then ρ conforms with σ_{Min}^0 , thus, as σ_{Min}^0 is fake-optimal, $\text{Price}_{\mathcal{G}_{L', r_0}}(\rho) \leq \text{Val}_{\mathcal{G}_{L', r_0}}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, \nu)$ (the last inequality comes from the definition of $r_1 = \text{left}_{L'}(r_0)$). Therefore, in the following cases, we assume that $\nu \in (r'', r_1)$.
- Consider then the case where ρ is of the form $(\ell, \nu) \xrightarrow{c} (\ell^f, \nu')$.

- if $\ell \in L' \cap L_{\text{Min}}$, ℓ is urgent in \mathcal{G}_{L',r_0} , thus $\nu' = \nu$. Furthermore, since ρ conforms with σ_{Min} , by construction of σ_{Min} , the choice of σ_{Min}^1 on (r'', r_1) consists in going to ℓ^f , thus, as observed above, $1 \wedge \neg 2$ holds for ℓ . Therefore,

$$\text{Val}_{\mathcal{G}}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, r_0) + \pi(\ell)(r_0 - \nu) = \varphi_{\ell^f}(\nu) = \text{Price}_{\mathcal{G}_{L',r_0}}(\rho).$$

- If $\ell \in L_{\text{Min}} \setminus L'$, by construction, it must be the case that $\sigma_{\text{Min}}(\ell, \nu) = (r_1 - \nu + t', (\ell, \ell^f))$ where $(t, (\ell, \ell^f)) = \sigma_{\text{Min}}^1(\ell, \nu)$ and $(t', (\ell, \ell^f)) = \sigma_{\text{Min}}^0(\ell, r_1)$. Thus, $\nu' = r_1 + t'$. In particular, observe that

$$\text{Price}_{\mathcal{G}_{L',r_0}}(\rho) = (r_1 - \nu)\pi(\ell) + \text{Price}_{\mathcal{G}_{L',r_0}}(\rho')$$

where $\rho' = (\ell, r_1) \xrightarrow{c'} (\ell^f, \nu')$. As ρ' conforms with σ_{Min}^0 which is fake-optimal in \mathcal{G}_{L',r_0} , and $r_1 = \text{left}_{L'}(r_0)$,

$$\text{Price}_{\mathcal{G}_{L',r_0}}(\rho') \leq \text{Val}_{\mathcal{G}_{L',r_0}}(\ell, r_1) = \text{Val}_{\mathcal{G}}(\ell, r_1).$$

Thus

$$\text{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leq (r_1 - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r_1) = \text{Price}_{\mathcal{G}_{L',r_1}}(\rho'')$$

where $\rho'' = (\ell, \nu) \xrightarrow{c''} (\ell^f, \nu + t)$ conforms with σ_{Min}^1 which is fake-optimal in \mathcal{G}_{L',r_1} . Therefore, since $r_1 = \text{left}_{L'}(r_0)$,

$$\text{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leq \text{Val}_{\mathcal{G}_{L',r_1}}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, \nu).$$

- If $\ell \in L_{\text{Max}}$ then

$$\begin{aligned} \text{Price}_{\mathcal{G}_{L',r_0}}(\rho) &= (\nu' - \nu)\pi(\ell) + \varphi_{\ell^f}(\nu') \\ &= (\nu' - \nu)\pi(\ell) + (r_0 - \nu')\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r_0) \\ &= (r_0 - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r_0). \end{aligned}$$

By Lemma 5.3, since $\ell \in L_{\text{Max}} \setminus (L_u \cup L_f)$ (ℓ is not urgent in \mathcal{G} since ℓ^f exists), $\text{Val}_{\mathcal{G}}(\ell, r_1) \geq (r_0 - r_1)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r_0)$. Furthermore, observe that if we define ρ' as the play $(\ell, \nu) \xrightarrow{c'} (\ell^f, \nu)$ in \mathcal{G}_{L',r_1} , then ρ' conforms with σ_{Min}^1 and

$$\begin{aligned} \text{Price}_{\mathcal{G}_{L',r_1}}(\rho') &= (r_1 - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r_1) \\ &\geq (r_1 - \nu)\pi(\ell) + (r_0 - r_1)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r_0) \\ &= (r_0 - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}}(\ell, r_0) \\ &= \text{Price}_{\mathcal{G}_{L',r_0}}(\rho). \end{aligned}$$

Thus, as σ_{Min}^1 is fake-optimal in \mathcal{G}_{L',r_1} ,

$$\text{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leq \text{Price}_{\mathcal{G}_{L',r_1}}(\rho') \leq \text{Val}_{\mathcal{G}_{L',r_1}}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, \nu).$$

- We finally consider the case where $\rho = (\ell, \nu) \xrightarrow{c} \rho'$ with ρ' that starts in configuration (ℓ', ν') such that $\ell' \notin \{\ell''^f \mid \ell'' \in L\}$. By induction hypothesis $\text{Price}_{\mathcal{G}_{L',r_0}}(\rho') \leq \text{Val}_{\mathcal{G}}(\ell', \nu')$.
 - If $\nu' \leq r_1$, let ρ'' be the play of \mathcal{G}_{L',r_1} starting in (ℓ', ν') that conforms with σ_{Min}^1 and σ_{Max}^1 . If ρ'' does not reach a final location, since σ_{Min}^1 is an NC-strategy, the costs of its prefixes tend to $-\infty$. By considering the switching strategy of Lemma 3.9, we would obtain a completed play conforming with σ_{Max}^1 of price smaller than $\text{Val}_{\mathcal{G}_{L',r_1}}(\ell', \nu')$ which would contradict the optimality of σ_{Max}^1 . Hence, ρ'' reaches the target. Moreover,

since σ_{Max}^1 is optimal and σ_{Min}^1 is fake-optimal, we finally know that $\text{Price}_{\mathcal{G}_{L',r_1}}(\rho'') = \text{Val}_{\mathcal{G}_{L',r_1}}(\ell', \nu')$ (since $\nu' \in [\text{left}_{L'}(r_1), r_1]$). Therefore,

$$\begin{aligned} \text{Price}_{\mathcal{G}_{L',r_0}}(\rho) &= (\nu' - \nu)\pi(\ell) + \pi(\ell, \ell') + \text{Price}_{\mathcal{G}_{L',r_0}}(\rho') \\ &\leq (\nu' - \nu)\pi(\ell) + \pi(\ell, \ell') + \text{Val}_{\mathcal{G}}(\ell', \nu') \\ &= (\nu' - \nu)\pi(\ell) + \pi(\ell, \ell') + \text{Price}(\rho'') = \text{Price}((\ell, \nu) \xrightarrow{c'} \rho'') \end{aligned}$$

Since the play $(\ell, \nu) \xrightarrow{c'} \rho''$ conforms with σ_{Min}^1 , we finally have

$$\text{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leq \text{Price}((\ell, \nu) \xrightarrow{c'} \rho'') \leq \text{Val}_{\mathcal{G}_{L',r_1}}(\ell, \nu) = \text{Val}_{\mathcal{G}}(\ell, \nu).$$

- If $\nu' > r_1$ and $\ell \in L_{\text{Max}}$, let ρ^1 be the play in \mathcal{G}_{L',r_1} defined by $\rho^1 = (\ell, \nu) \xrightarrow{c'} (\ell^f, \nu)$ and ρ^0 the play in \mathcal{G}_{L',r_0} defined by $\rho^0 = (\ell, r_1) \xrightarrow{c''} \rho'$. We have

$$\begin{aligned} \text{Price}_{\mathcal{G}_{L',r_0}}(\rho) &= (\nu' - \nu)\pi(\ell) + \pi(\ell, \ell') + \text{Price}_{\mathcal{G}_{L',r_0}}(\rho') \\ &= \underbrace{\varphi_{\ell^f}(\nu) - \text{Val}_{\mathcal{G}}(\ell, r_1)}_{= \text{Price}_{\mathcal{G}_{L',r_1}}(\rho^1)} + \underbrace{(\nu' - r_1)\pi(\ell) + \pi(\ell, \ell') + \text{Price}_{\mathcal{G}_{L',r_0}}(\rho')}_{= \text{Price}_{\mathcal{G}_{L',r_0}}(\rho^0)}. \end{aligned}$$

Since ρ^0 conforms with σ_{Min}^0 , fake-optimal, and reaches a final location, and since $r_1 = \text{left}_{L'}(r_0)$,

$$\text{Price}_{\mathcal{G}_{L',r_0}}(\rho^0) \leq \text{Val}_{\mathcal{G}_{L',r_0}}(\ell, r_1) = \text{Val}_{\mathcal{G}}(\ell, r_1).$$

We also have that ρ^1 conforms with σ_{Min}^1 , so the previous explanations already proved that $\text{Price}_{\mathcal{G}_{L',r_1}}(\rho^1) \leq \text{Val}_{\mathcal{G}}(\ell, \nu)$. As a consequence $\text{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leq \text{Val}_{\mathcal{G}}(\ell, \nu)$.

- If $\nu' > r_1$ and $\ell \in L_{\text{Min}}$, we know that ℓ is non-urgent, so that $\ell \notin L'$. Therefore, by definition of σ_{Min} , $\sigma_{\text{Min}}(\ell, \nu) = (r_1 - \nu + t', (\ell, \ell'))$ where $\sigma_{\text{Min}}^1(\ell, \nu) = (t, (\ell, \ell^f))$ for some delay t , and $\sigma_{\text{Min}}^0(\ell, r_1) = (t', (\ell, \ell'))$. If we let ρ^1 be the play in \mathcal{G}_{L',r_1} defined by $\rho^1 = (\ell, \nu) \xrightarrow{c'} (\ell^f, \nu)$ and ρ^0 the play in \mathcal{G}_{L',r_0} defined by $\rho^0 = (\ell, r_1) \xrightarrow{c''} \rho'$, as in the previous case, we obtain that $\text{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leq \text{Val}_{\mathcal{G}}(\ell, \nu)$.

As a consequence of this induction, we have shown that for all $\ell \in L$, and $\nu \in (r'', r_1)$, $\text{fake}_{\mathcal{G}_{L',r_0}}^{\sigma_{\text{Min}}}(\ell, \nu) \leq \text{Val}_{\mathcal{G}}(\ell, \nu)$, which shows one inequality of (5.4), the other being obtained very similarly. \square

Next, we will define two different ways of choosing the subset $L' \subseteq L \setminus (L_u \cup L_f)$: the former (one at a time) to prove finite optimality of all SPTGs, the latter (all at once) to bound the number of cutpoints of the value functions and obtain an efficient algorithm to solve them.

5.2. SPTGs are finitely optimal. To prove finite optimality of all SPTGs we reason by induction on the number of non-urgent locations and instantiate the previous results to the case where $L' = \{\ell^*\}$ where ℓ^* is a non-urgent location of *minimum weight* (i.e. for all $\ell \in L \setminus (L_f \cup L_u)$, $\pi(\ell^*) \leq \pi(\ell)$). Given $r_0 \in [0, 1]$, we let $r_0 > r_1 > \dots$ be the decreasing sequence of clock values such that $r_i = \text{left}_{\ell^*}(r_{i-1})$ for all $i > 0$ with $r_{i-1} > 0$. As explained before, we will build $\text{Val}_{\mathcal{G}}$ on $[\inf_i r_i, r_0]$ from the value functions of games $\mathcal{G}_{\ell^*, r_i}$. Assuming finite optimality of those games, this will prove that \mathcal{G} is finitely optimal *under the condition*




Figure 7: In this example $L' = \{\ell^*\}$ and $\ell^* \in L_{\text{Min}}$. $\text{left}_{\ell^*}(r)$ is the leftmost point such that all slopes on its right are at least $-\pi(\ell^*)$ in the graph of $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu)$. Dashed lines have slope $-\pi(\ell^*)$.

that $r_0 > r_1 > \dots$ eventually stops, i.e. $r_i = 0$ for some i . Lemma 5.8 will prove this property. First, we relate the optimal value functions with the final cost functions.

Lemma 5.7. *Assume that $\mathcal{G}_{\ell^*, r}$ is finitely optimal. If $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*)$ is affine on a non-singleton interval $I \subseteq [0, r]$ with a slope greater⁷ than $-\pi(\ell^*)$, then there exists $f \in F_{\mathcal{G}}$ (see definition in page 21) such that for all $\nu \in I$, $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu) = f(\nu)$.*

Proof. Let σ_{Min}^1 and σ_{Max} be some fake-optimal NC-strategy and optimal FP-strategy in $\mathcal{G}_{\ell^*, r}$. As I is a non-singleton interval, there exists a subinterval $I' \subset I$, which is not a singleton and is contained in an interval of σ_{Min}^1 and of σ_{Max} . Let σ_{Min} be the switching strategy obtained from σ_{Min}^1 in Lemma 3.9: notice that both strategies have the same intervals.

Let $\nu \in I'$. Since $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu) \notin \{+\infty, -\infty\}$, the completed play

$$\text{CPlay}((\ell^*, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})$$

necessarily reaches a final location and has price $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu)$. Thus it is a finite completed play $(\ell_0, \nu_0) \xrightarrow{c_0} \dots (\ell_k, \nu_k)$ where $(\ell_0, \nu_0) = (\ell^*, \nu)$ and $\ell_k \in L'_f$. We also let $\nu' \in I'$ be a clock value such that $\nu < \nu'$. We now explain successively why:

- (1) for all i , $\nu_i = \nu$;
- (2) $\ell_k \in L'_f$;
- (3) $\text{CPlay}((\ell^*, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})$ contains no cycles.

We will then use these properties to conclude.

- (1) Assume by contradiction that there exists an index i such that $\nu < \nu_i$ and let i be the smallest of such indices. For each $j < i$, if $\ell_j \in L_{\text{Min}}$, let $(t, \delta) = \sigma_{\text{Min}}(\ell_j, \nu)$ and $(t', \delta') = \sigma_{\text{Min}}(\ell_j, \nu')$. Similarly, if $\ell_j \in L_{\text{Max}}$, we let $(t, \delta) = \sigma_{\text{Max}}(\ell_j, \nu)$ and $(t', \delta') = \sigma_{\text{Max}}(\ell_j, \nu')$. As I' is contained in an interval of σ_{Min} and σ_{Max} , we have $\delta = \delta'$

⁷For this result, the order does not depend on the owner of the location, but on the fact that ℓ^* has minimal weight amongst locations of \mathcal{G} .

and either $t = t' = 0$, or $\nu + t = \nu' + t'$. Applying this result for all $j < i$, we obtain that $(\ell_0, \nu') \xrightarrow{c'_0} \cdots (\ell_{i-1}, \nu') \xrightarrow{c'_{i-1}} (\ell_i, \nu_i) \xrightarrow{c_i} \cdots (\ell_k, \nu_k)$ is a prefix of $\text{CPlay}((\ell^*, \nu'), \sigma_{\text{Min}}, \sigma_{\text{Max}})$: notice moreover that, as before, this prefix has cost $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu')$. In particular,

$$\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu') = \text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu) - (\nu' - \nu)\pi(\ell_{i-1}) \leq \text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu) - (\nu' - \nu)\pi(\ell^*)$$

which implies that the slope of $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*)$ is at most $-\pi(\ell^*)$, and therefore contradicts the hypothesis. As a consequence, we have that $\nu_i = \nu$ for all i .

- (2) Again by contradiction, assume now that $\ell_k = \ell^f$ for some $\ell \in L \setminus (L_u \cup L_f)$. By the same reasoning as before, we then would have $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu') = \text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu) - (\nu' - \nu)\pi(\ell)$, which again contradicts the hypothesis. Therefore, $\ell_k \in L_f$.
- (3) Suppose, for a contradiction, that the prefix $(\ell_0, \nu) \xrightarrow{c_0} \cdots (\ell_k, \nu)$ contains a cycle. Since σ_{Min} is a switching strategy and σ_{Max} is a memoryless strategy, this implies that the cycle is contained in the part of σ_{Min} where the decision is taken by the strategy σ_{Min}^1 : since it is an NC-strategy, this implies that the sum of the weights along the cycle is at most -1 . But if this is the case, we may modify the switching strategy σ_{Min} to loop more in the same cycle (this is indeed a cycle in the timed game, not only in the untimed region game): against the optimal memoryless strategy σ_{Max} , this would imply that Min has a sequence of strategies to obtain a value as small as it wants, and thus $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu) = -\infty$. This contradicts the absence of values $-\infty$ in the game. Thus, the prefix $(\ell_0, \nu) \xrightarrow{c_0} \cdots (\ell_k, \nu)$ contains no cycles.

We now explain how to conclude. The absence of cycles implies that the sum of the discrete weights $w = \pi(\ell_0, \ell_1) + \cdots + \pi(\ell_{k-1}, \ell_k)$ belongs to the set $[-(|L| - 1)\Pi^{\text{tr}}, (|L| - 1)\Pi^{\text{tr}}] \cap \mathbb{Z}$, and we have $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu) = w + \varphi_{\ell_k}(\nu)$. Notice that the previous developments also show that for all $\nu' \in I'$ (here, $\nu < \nu'$ is not needed), $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*, \nu') = w + \varphi_{\ell_k}(\nu')$, with the same location ℓ_k , and length k . Since this equality holds on $I' \subseteq I$ which is not a singleton, and $\text{Val}_{\mathcal{G}_{\ell^*, r}}(\ell^*)$ is affine on I , it holds everywhere on I . This shows the result since $w + \varphi_{\ell_k} \in \text{FG}$. \square

We now prove the termination of the sequence of r_i 's described earlier. This is achieved by showing why, for all i , the owner of ℓ^* has a strictly better strategy in configuration (ℓ^*, r_{i+1}) than waiting until r_i in location ℓ^* .

Lemma 5.8. *If $\mathcal{G}_{\ell^*, r_i}$ is finitely optimal for all $i \geq 0$ for which r_i is defined, then*

- (1) *there exists $j \leq |\text{FG}|^2 + 2$ such that $r_j = 0$; and*
- (2) *denoting j the number such that $r_j = 0$ we have for all $0 \leq i \leq j - 2$ that if $\ell^* \in L_{\text{Min}}$ (respectively, L_{Max}), $\text{Val}_{\mathcal{G}}(\ell^*, r_{i+1}) < \text{Val}_{\mathcal{G}}(\ell^*, r_i) + (r_i - r_{i+1})\pi(\ell^*)$ (respectively, $\text{Val}_{\mathcal{G}}(\ell^*, r_{i+1}) > \text{Val}_{\mathcal{G}}(\ell^*, r_i) + (r_i - r_{i+1})\pi(\ell^*)$).*

Proof. (1) We consider first the case where $\ell^* \in L_{\text{Max}}$, showing a better bound $j \leq |\text{FG}| + 2$. The main ingredient is to show that a function of FG cannot be used twice in $\text{Val}_{\mathcal{G}}(\ell^*)$. Let $i > 0$ such that $r_i \neq 0$ (if there exist no such i then $r_1 = 0$). Recall from Lemma 5.5 that there exists $r'_i < r_i$ such that $\text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*)$ is affine on $[r'_i, r_i]$, of slope greater than $-\pi(\ell^*)$. In particular,

$$\frac{\text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*, r_i) - \text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*, r'_i)}{r_i - r'_i} > -\pi(\ell^*).$$




Figure 8: The case $\ell^* \in L_{\text{Max}}$: a geometric proof of $f_i \neq f_j$. The dotted lines represents f_i and f_j , the dashed lines have slope $-\pi(\ell^*)$, and the plain line depicts $\text{Val}_{\mathcal{G}}(\ell^*)$. Because the slope of f_i is strictly smaller than $-\pi(\ell^*)$, and the value at r_j is above the dashed line it cannot be the case that $f_i(r_j) = \text{Val}_{\mathcal{G}}(\ell^*, r_j) = f_j(r_j)$.

Lemma 5.7 states that on $[r'_i, r_i]$, $\text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*)$ is equal to some $f_i \in \mathbb{F}_{\mathcal{G}}$. As f_i is an affine function, $f_i(r_i) = \text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*, r_i)$, and $f_i(r'_i) = \text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*, r'_i)$. Thus, for all ν ,

$$f_i(\nu) = \text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*, r_i) + \frac{\text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*, r'_i) - \text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*, r_i)}{r_i - r'_i}(r_i - \nu).$$

Since $\mathcal{G}_{\ell^*, r_{i-1}}$ is assumed to be finitely optimal, we know that $\text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}(\ell^*, r_i) = \text{Val}_{\mathcal{G}}(\ell^*, r_i)$, by definition of $r_i = \text{left}_{\ell^*}(r_{i-1})$. Therefore, combining both equalities above, for all clock values $\nu < r_i$, we have $f_i(\nu) < \text{Val}_{\mathcal{G}}(\ell^*, r_i) + \pi(\ell^*)(r_i - \nu)$.

Consider then $j > i$ such that $r_j \neq 0$. We claim that $f_j \neq f_i$. Indeed, we have $\text{Val}_{\mathcal{G}}(\ell^*, r_j) = f_j(r_j)$. As, in \mathcal{G} , ℓ^* is a non-urgent location, Lemma 5.3 ensures that

$$\text{Val}_{\mathcal{G}}(\ell^*, r_j) \geq \text{Val}_{\mathcal{G}}(\ell^*, r_i) + \pi(\ell^*)(r_i - r_j).$$

As for all i' , $\text{Val}_{\mathcal{G}}(\ell^*, r_{i'}) = f_{i'}(r_{i'})$, the equality above is equivalent to $f_j(r_j) \geq f_i(r_i) + \pi(\ell^*)(r_i - r_j)$. Recall that f_i has a slope strictly greater than $-\pi(\ell^*)$, therefore $f_i(r_j) < f_i(r_i) + \pi(\ell^*)(r_i - r_j) \leq f_j(r_j)$. As a consequence $f_i \neq f_j$ (this is depicted in Figure 8).

Therefore, there cannot be more than $|\mathbb{F}_{\mathcal{G}}| + 1$ non-null elements in the sequence $r_0 \geq r_1 \geq \dots$, which proves that there exists $i \leq |\mathbb{F}_{\mathcal{G}}| + 2$ such that $r_i = 0$.

We continue with the case where $\ell^* \in L_{\text{Min}}$. We generalise the previous arguments that may no longer be true in this case (the same function of $\mathbb{F}_{\mathcal{G}}$ could be used twice in $\text{Val}_{\mathcal{G}}(\ell^*)$), by showing that in-between two successive points r_{i+1} and r_i , there is always one “full segment” of $\mathbb{F}_{\mathcal{G}}$ (i.e. it encounters at least one point that is the intersection of two functions of $\mathbb{F}_{\mathcal{G}}$, and there are $|\mathbb{F}_{\mathcal{G}}|^2$ many such points). Let $r_{\infty} = \inf\{r_i \mid i \geq 0\}$. In this case, we look at the affine parts of $\text{Val}_{\mathcal{G}}(\ell^*)$ with a slope greater than $-\pi(\ell^*)$, and we show that there can only be finitely many such segments in $[r_{\infty}, 1]$. We then show that there is at least one such segment contained in $[r_{i+1}, r_i]$ for all i , bounding the size of the sequence.




Figure 9: In order for the segments $[a, b]$ and $[c, d]$ to be aligned, there must exist a segment with a biggest slope crossing $f_{[a,b]}$ (represented by a dashed line) between b and c .

In the following, we call *segment* every interval $[a, b] \subset (r_\infty, 1]$ such that a and b are two consecutive cutpoints of the cost function $\text{Val}_G(\ell^*)$ over $(r_\infty, 1]$. Recall that it means that $\text{Val}_G(\ell^*)$ is affine on $[a, b]$, and if we let a' be the greatest cutpoint smaller than a , and b' be the smallest cutpoint greater than b , the slopes of $\text{Val}_G(\ell^*)$ on $[a', a]$ and $[b, b']$ are different from the slope on $[a, b]$. We abuse the notations by referring to *the slope of a segment* $[a, b]$ for the slope of $\text{Val}_G(\ell^*)$ on $[a, b]$ and simply call *cutpoint* a cutpoint of $\text{Val}_G(\ell^*)$.

To every segment $[a, b]$ with a slope greater than $-\pi(\ell^*)$, we associate a function $f_{[a,b]} \in \mathcal{F}_G$ as follows. Let i be the smallest index such that $[a, b] \cap [r_{i+1}, r_i]$ is a non singleton interval $[a', b']$. Lemma 5.7 ensures that there exists $f_{[a,b]} \in \mathcal{F}_G$ such that for all $\nu \in [a', b']$, $\text{Val}_G(\ell^*, \nu) = f_{[a,b]}(\nu)$.

Consider now two disjoint segments $[a, b]$ and $[c, d]$ with a slope greater than $-\pi(\ell^*)$, and assume that $f_{[a,b]} = f_{[c,d]}$ (in particular both segments have the same slope). Without loss of generality, assume that $b < c$. We claim that there exists a segment $[e, g]$ in-between $[a, b]$ and $[c, d]$ with a slope greater than the slope of $[c, d]$, and that $f_{[e,g]}$ and $f_{[a,b]}$ intersect over $x \in [b, c]$, i.e. $f_{[e,g]}(x) = f_{[a,b]}(x)$ (depicted in Figure 9). We prove it now.

Let α be the greatest cutpoint smaller than c . We know that the slope of $[\alpha, c]$ is different from the one of $[c, d]$. If it is greater then define $e = \alpha$ and $x = g = c$, those indeed satisfy the property. If the slope of $[\alpha, c]$ is smaller than the one of $[c, d]$, then for all $\nu \in [\alpha, c)$, $\text{Val}_G(\ell^*, \nu) > f_{[c,d]}(\nu)$. Let x be the greatest point in $[b, \alpha]$ such that $\text{Val}_G(\ell^*, x) = f_{[c,d]}(x)$. We know that it exists since $\text{Val}_G(\ell^*, b) = f_{[c,d]}(b)$, and $\text{Val}_G(\ell^*)$ is continuous. Observe that $\text{Val}_G(\ell^*, \nu) > f_{[c,d]}(\nu)$, for all $x < \nu < c$. Finally, let g be the smallest cutpoint of $\text{Val}_G(\ell^*)$ strictly greater than x , and e the greatest cutpoint of $\text{Val}_G(\ell^*)$ smaller than or equal to x . By construction, $[e, g]$ is a segment that contains x . The slope of the segment $[e, g]$ is $s_{[e,g]} = \frac{\text{Val}_G(\ell^*, g) - \text{Val}_G(\ell^*, x)}{g - x}$,




Figure 10: The case $\ell^* \in L_{\text{Min}}$: as the value at r_{i+1} is strictly below $\text{Val}_{\mathcal{G}}(\ell^*, r_i) + \pi(\ell^*)(r_i - r_{i+1})$, as the slope on the left of r_i and of r_{i+1} is $-\pi(\ell^*)$, there must exist a segment (represented with a double line) with slope greater than $-\pi(\ell^*)$ in $[r_{i+1}, r_i]$.

and the slope of the segment $[c, d]$ is equal to $s_{[c,d]} = \frac{f_{[c,d]}(g) - f_{[c,d]}(x)}{g - x}$. Remembering that $\text{Val}_{\mathcal{G}}(\ell^*, x) = f_{[c,d]}(x)$, and that $\text{Val}_{\mathcal{G}}(\ell^*, g) > f_{[c,d]}(g)$ since $g \in (x, c)$, we obtain that $s_{[e,g]} > s_{[c,d]}$. Finally, since $\text{Val}_{\mathcal{G}}(\ell^*, x) = f_{[c,d]}(x) = f_{[e,g]}(x)$, x is indeed the intersection point of $f_{[c,d]} = f_{[a,b]}$ and $f_{[e,g]}$, which concludes the proof of the previous claim.

For every function $f \in \mathcal{F}_{\mathcal{G}}$, there are less than $|\mathcal{F}_{\mathcal{G}}|$ intersection points between f and the other functions of $\mathcal{F}_{\mathcal{G}}$ (at most one for each pair (f, f')). If f has a slope greater than $-\pi(\ell^*)$, thanks to the previous paragraph, we know that there are at most $|\mathcal{F}_{\mathcal{G}}|$ segments $[a, b]$ such that $f_{[a,b]} = f$. Summing over all possible functions f , there are at most $|\mathcal{F}_{\mathcal{G}}|^2$ segments with a slope greater than $-\pi(\ell^*)$.

Now, we link those segments with the clock values r_i 's, for $i > 0$. By item 2, thanks to the finite-optimality of $\mathcal{G}_{\ell^*, r_i}$, $\text{Val}_{\mathcal{G}}(\ell^*, r_{i+1}) < (r_i - r_{i+1})\pi(\ell^*) + \text{Val}_{\mathcal{G}}(\ell^*, r_i)$. Furthermore, Lemma 5.6 states that the slope of the segment directly on the left of r_i is equal to $-\pi(\ell^*)$. With the previous inequality in mind, this cannot be the case if $\text{Val}_{\mathcal{G}}(\ell^*)$ is affine over the whole interval $[r_{i+1}, r_i]$. Thus, there exists a segment $[a, b]$ of slope strictly greater than $-\pi(\ell^*)$ such that $b \in [r_{i+1}, r_i]$. As we also know that the slope left to r_{i+1} is $-\pi(\ell^*)$, it must be the case that $a \in [r_{i+1}, r_i]$. Hence, we have shown that in-between r_{i+1} and r_i , there is always a segment (this is depicted in Figure 10). As the number of such segments is bounded by $|\mathcal{F}_{\mathcal{G}}|^2$, we know that the sequence r_i is stationary in at most $|\mathcal{F}_{\mathcal{G}}|^2 + 1$ steps, i.e. that there exists $i \leq |\mathcal{F}_{\mathcal{G}}|^2 + 1$ such that $r_i = 0$.

- (2) We assume $\ell^* \in L_{\text{Min}}$, since the proof of the other case only differs with respect to the sense of the inequalities. From Lemma 5.5, we know that in $\mathcal{G}_{\ell^*, r_i}$, if $r_{i+1} > 0$, there exists $r' < r_{i+1}$ such that $\text{Val}_{\mathcal{G}_{\ell^*, r_i}}(\ell^*)$ is affine on $[r', r_{i+1}]$ and its slope is smaller than $-\pi(\ell^*)$, i.e. $\text{Val}_{\mathcal{G}_{\ell^*, r_i}}(\ell^*, r_{i+1}) < \text{Val}_{\mathcal{G}_{\ell^*, r_i}}(\ell^*, r') - (r_{i+1} - r')\pi(\ell^*)$. Lemma 5.4 also ensures that $\text{Val}_{\mathcal{G}_{\ell^*, r_i}}(\ell^*, r') \leq \text{Val}_{\mathcal{G}}(\ell^*, r_i) + (r_i - r')\pi(\ell^*)$. Combining both inequalities allows us to conclude. \square

We iterate this construction to obtain the finite optimality:

Theorem 5.9. *Every SPTG \mathcal{G} is finitely optimal.*

Proof. We show by induction on $n \geq 0$ that every r -SPTG \mathcal{G} with n non-urgent non-final locations is finitely optimal.

The base case $n = 0$ is given by Proposition 4.7.

Now, assume that \mathcal{G} has at least one non-urgent location, and assume ℓ^* is a non-urgent location with minimum weight. By induction hypothesis, all r' -SPTGs $\mathcal{G}_{\ell^*, r'}$ are finitely optimal for all $r' \in [0, r]$. Let $r_0 > r_1 > \dots$ be the decreasing sequence defined by $r_0 = r$ and $r_i = \text{left}_{\ell^*}(r_{i-1})$ for all $i \geq 1$. By Lemma 5.8, there exists $j \leq |\mathbb{F}_{\mathcal{G}}|^2 + 2$ such that $r_j = 0$. Moreover, for all $0 < i \leq j$, $\text{Val}_{\mathcal{G}} = \text{Val}_{\mathcal{G}_{\ell^*, r_{i-1}}}$ on $[r_i, r_{i-1}]$ by definition of $r_i = \text{left}_{\ell^*}(r_{i-1})$, so that $\text{Val}_{\mathcal{G}}(\ell)$ is a cost function on this interval by induction hypothesis. Finally, by using Proposition 5.2, we can reconstruct fake-optimal and optimal strategies in \mathcal{G} from the fake-optimal and optimal strategies of $\mathcal{G}_{\ell^*, r_i}$. \square

5.3. SPTGs have a pseudo-polynomial number of cutpoints. To prove that the number of cutpoints of value functions of SPTGs is at most pseudo-polynomial, we will need more knowledge about the left_{ℓ^*} operator for all SPTGs \mathcal{G} . First, if we are at the left of a given position r_1 , the next jump is further than $\text{left}_{\ell^*}(r_1)$:

Lemma 5.10. *Let \mathcal{G} be an SPTG and ℓ^* be a non-urgent location of minimum weight. Let r_1 and r_2 be such that $r_2 \leq r_1$. Then, $\text{left}_{\ell^*}(r_2) \leq \text{left}_{\ell^*}(r_1)$.*

Proof. Let $r'_1 = \text{left}_{\ell^*}(r_1)$. If $r_2 \leq r'_1$, then the result is trivially true. We now suppose that $r'_1 < r_2 < r_1$. By definition, it suffices to show that

$$\forall \nu \in [r'_1, r_2] \quad \text{Val}_{\mathcal{G}_{\ell^*, r_1}}(\ell^*, \nu) = \text{Val}_{\mathcal{G}_{\ell^*, r_2}}(\ell^*, \nu) \quad (5.5)$$

Indeed, since $\text{Val}_{\mathcal{G}_{\ell^*, r_1}}(\ell^*, \nu) = \text{Val}_{\mathcal{G}}(\ell^*, \nu)$ for all $\nu \in [r'_1, r_1]$, this implies that for all $\nu \in [r'_1, r_2]$, $\text{Val}_{\mathcal{G}_{\ell^*, r_2}}(\ell^*, \nu) = \text{Val}_{\mathcal{G}}(\ell^*, \nu)$, and thus that $\text{left}_{\ell^*}(r_2) \leq r'_1$.

To show (5.5), for $\varepsilon > 0$, we will need ε -optimal strategies for both players in $\mathcal{G}_{\ell^*, r_1}$ and \mathcal{G}_{r_1} . We build them considering two separate cases.

- if ℓ^* belongs to Max, let σ_{Max}^* be an ε -optimal strategy of Max in $\mathcal{G}_{\ell^*, r_1}$: then it is also ε -optimal strategy in \mathcal{G}_{r_1} since the values of those games are the same on $[r'_1, r_1]$, and strategies of Min are identical too;
- if ℓ^* belongs to Min, let σ_{Max}^* be an ε -optimal strategy of Max in \mathcal{G}_{r_1} : then it is also ε -optimal in $\mathcal{G}_{\ell^*, r_1}$ since Min has less capabilities in this game than in \mathcal{G}_{r_1} , while strategies of Max are unchanged.

We do the same case distinction to define an ε -optimal strategy σ_{Min}^* both in $\mathcal{G}_{\ell^*, r_1}$ and \mathcal{G}_{r_1} . In particular, we have, for all $\nu \in [r'_1, r_1]$

$$\text{Val}_{\mathcal{G}_{r_1}}(\ell, \nu) - \varepsilon \leq \text{Val}_{\mathcal{G}_{r_1}}^{\sigma_{\text{Max}}^*}(\ell, \nu) \quad \text{and} \quad \text{Val}_{\mathcal{G}_{r_1}}^{\sigma_{\text{Min}}^*}(\ell, \nu) \leq \text{Val}_{\mathcal{G}_{r_1}}(\ell, \nu) + \varepsilon$$

First, let us show that $\text{Val}_{\mathcal{G}_{\ell^*, r_1}}(\ell^*, \nu) \leq \text{Val}_{\mathcal{G}_{\ell^*, r_2}}(\ell^*, \nu)$, i.e. $\text{Val}_{\mathcal{G}_{r_1}}(\ell^*, \nu) \leq \text{Val}_{\mathcal{G}_{\ell^*, r_2}}(\ell^*, \nu)$. To do so, we consider the definition of the value and show that for all $\nu \in [r'_1, r_2]$,

$$\text{Val}_{\mathcal{G}_{r_1}}^{\sigma_{\text{Max}}^*}(\ell^*, \nu) \leq \sup_{\sigma_{\text{Max}}} \inf_{\sigma_{\text{Min}}} \text{Price}(\text{CPlay}((\ell^*, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) + \varepsilon \quad (5.6)$$

where the play on the right is a play of the game $\mathcal{G}_{\ell^*, r_2}$. In particular, this will imply that $\text{Val}_{\mathcal{G}_{r_1}}(\ell^*, \nu) - \varepsilon \leq \text{Val}_{\mathcal{G}_{\ell^*, r_2}}(\ell^*, \nu) + \varepsilon$ which allows us to conclude by letting ε go to 0. We

thus build a strategy σ_{Max} in $\mathcal{G}_{\ell^*, r_2}$ as follows: it simply follows what σ_{Max}^* prescribes to do in \mathcal{G}_{r_1} (especially when it jumps from ℓ^* to ℓ^{*f} if ℓ^* belongs to Max) except when σ_{Max}^* wants to jump on the right of r_2 , from any location ℓ , in which case σ_{Max} goes to the location ℓ^f in valuation r_2 . We now explain why

$$\text{Val}_{\mathcal{G}_{r_1}}^{\sigma_{\text{Max}}^*}(\ell^*, \nu) \leq \inf_{\sigma_{\text{Min}}} \text{Price}(\text{CPlay}((\ell^*, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) + \varepsilon$$

To do so, we consider any strategy σ_{Min} of Min in $\mathcal{G}_{\ell^*, r_2}$, and build a strategy σ'_{Min} of Min in \mathcal{G}_{r_1} that gets a smaller payoff. The strategy σ'_{Min} mimics σ_{Min} except when it jumps in a location ℓ^f : instead σ'_{Min} delays in ℓ until r_2 and then performs the action prescribed by σ_{Min}^* in (ℓ, r_2) . Notice that this is a legal move since we play in \mathcal{G}_{r_1} where the location ℓ^* has not been made urgent.

We now compare the prices of two plays:

- the play ρ_1 obtained from (ℓ, ν) in \mathcal{G}_{r_1} by following σ_{Max}^* (the ε -optimal strategy we have fixed) and σ'_{Min} (that we have built);
- and the play ρ_2 obtained from (ℓ, ν) in $\mathcal{G}_{\ell^*, r_2}$ by following σ_{Max} (that we have built) and σ_{Min} (that we have fixed).

We need to show that $\text{Price}(\rho_1) \leq \text{Price}(\rho_2) + \varepsilon$ to conclude.

If ρ_2 stops in a location different from ℓ^f for any ℓ , then this play is also a play of \mathcal{G}_{r_1} conforming to σ_{Max}^* (by construction of σ_{Max}) and σ'_{Min} (that we have built), and is thus equal to ρ_1 . We conclude directly that $\text{Price}(\rho_1) = \text{Price}(\rho_2)$.

Otherwise, ρ_2 stops in a configuration (ℓ^f, ν) . Let ρ'_2 be the partial play obtained from ρ_2 by removing its last transition. Then,

$$\text{Price}(\rho_2) = \text{Price}(\rho'_2) + (r_2 - \nu)\pi(\ell) + \text{Val}_{\mathcal{G}_{\ell^*, r_1}}(\ell, r_2)$$

Let ρ'_1 be the play obtained by following σ_{Min}^* and σ_{Max}^* in \mathcal{G}_{r_1} from (ℓ, r_2) : ρ'_1 has price at most $\text{Val}_{\mathcal{G}_{r_1}}^{\sigma_{\text{Min}}^*}(\ell, r_2) \leq \text{Val}_{\mathcal{G}_{r_1}}(\ell, r_2) + \varepsilon$ since it follows σ_{Min}^* . However, the play ρ_1 is the concatenation of the play ρ'_2 , a delay of $r_2 - \nu$ in ℓ , and the play ρ'_1 . Thus

$$\begin{aligned} \text{Price}(\rho_1) &= \text{Price}(\rho'_2) + (r_2 - \nu)\pi(\ell) + \text{Price}(\rho'_1) \\ &= \text{Price}(\rho_2) - \text{Val}_{\mathcal{G}_{\ell^*, r_1}}(\ell, r_2) + \text{Price}(\rho'_1) \\ &\leq \text{Price}(\rho_2) + \varepsilon \end{aligned}$$

This concludes all the cases and thus the proof of (5.6).

The other inequality $\text{Val}_{\mathcal{G}_{\ell^*, r_1}}(\ell^*, \nu) \geq \text{Val}_{\mathcal{G}_{\ell^*, r_2}}(\ell^*, \nu)$ is obtained symmetrically by showing that

$$\text{Val}_{\mathcal{G}_{r_1}}^{\sigma_{\text{Min}}^*}(\ell^*, \nu) \geq \inf_{\sigma_{\text{Min}}} \sup_{\sigma_{\text{Max}}} \text{Price}(\text{CPlay}((\ell^*, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})) - \varepsilon \quad \square$$

Then, we change our policy to make locations urgent: instead of making them urgent one by one by increasing order of weight, we make them all urgent at once. We now show that this makes us progress at least as fast in the left functions (from now on, we reuse the exponents in the left function to explain which game we consider):

Lemma 5.11. *Let \mathcal{G} be an SPTG with non-urgent locations $L' = \{\ell_1, \dots, \ell_n\}$ ordered in increasing order of weight. Then for all valuations $r \leq 1$,*

$$\text{left}_{L'}^{\mathcal{G}}(r) \leq \max_{1 \leq i \leq n} \text{left}_{\ell_i}^{\mathcal{G}_{\{\ell_1, \dots, \ell_{i-1}\}, r}}(r)$$

Proof. Let $r' = \max_{1 \leq i \leq n} \text{left}_{\ell_i}^{\mathcal{G}_{\{\ell_1, \dots, \ell_{i-1}\}, r}}(r)$. By definition of left function, for all $\ell \in L$, $0 \leq i \leq n-1$, and $\nu \in [r', r]$, $\text{Val}_{\mathcal{G}_{\ell_0, \dots, \ell_{i-1}}}(\ell, \nu) = \text{Val}_{\mathcal{G}_{\ell_0, \dots, \ell_i}}(\ell, \nu)$. Thus $\text{Val}_{\mathcal{G}}(\ell, \nu) = \text{Val}_{\mathcal{G}_{L', r}}(\ell, \nu)$ for all $\nu \in [r', r]$, meaning that $\text{left}_{L'}^{\mathcal{G}}(r) \leq r'$. \square

This allows us to bound the number of steps of left when making all locations urgent at once, generalising the bound obtained in Lemma 5.8 when making only one location urgent:

Lemma 5.12. *Let \mathcal{G} be an SPTG with non-urgent locations L' . Denoting $(r_k)_{k \in \mathbb{N}}$ the sequence defined by $r_0 = 1$ and for all i , $r_{i+1} = \text{left}_{L'}^{\mathcal{G}}(r_i)$, then there exists $j \leq |L|(|\mathbb{F}_{\mathcal{G}}|^2 + 2)$ such that $r_j = 0$.*

Proof. Let ℓ_1, \dots, ℓ_n be the locations of L' by increasing order of weight. By Lemma 5.8, for all $1 \leq i \leq n$, the sequence $(j_k^{(i)})_{k \in \mathbb{N}}$ defined by $j_0^{(i)} = 1$ and $j_{k+1}^{(i)} = \text{left}_{\ell_i}^{\mathcal{G}_{\ell_1, \dots, \ell_{i-1}}}(j_k^{(i)})$ for all $k \in \mathbb{N}$ is stationary to 0: there exists $k_i \leq |\mathbb{F}_{\mathcal{G}}|^2 + 2$ such that $j_{k_i}^{(i)} = 0$. We now build the decreasing sequence $(j_k)_{0 \leq k \leq t}$ by interleaving those n sequences. We thus have $t \leq |L'|(|\mathbb{F}_{\mathcal{G}}|^2 + 2) \leq |L|(|\mathbb{F}_{\mathcal{G}}|^2 + 2)$.

For all $k \leq t$, we have $r_k \leq j_k$. Indeed, $r_0 = 1 = j_0$ (as $j_0^{(i)} = 1$ for all $i \leq n$). Assume that $r_k \leq j_k$ for some $k \leq t$ and, for all $1 \leq i \leq n$, let n_i be the greatest index such that $j_k \leq j_{n_i}^{(i)}$, so that $r_k \leq j_{n_i}^{(i)}$. By definition of the sequence $(j_k)_{0 \leq k \leq t}$, we then have $j_{k+1} = \max_{1 \leq i \leq n} (j_{n_i+1}^{(i)})$. Thus

$$\begin{aligned} r_{k+1} = \text{left}_{L'}^{\mathcal{G}}(r_k) &\leq \max_{1 \leq i \leq n} \text{left}_{\ell_i}^{\mathcal{G}_{\ell_1, \dots, \ell_{i-1}}}(r_k) && \text{(by Lemma 5.11)} \\ &\leq \max_{1 \leq i \leq n} \text{left}_{\ell_i}^{\mathcal{G}_{\ell_1, \dots, \ell_{i-1}}}(j_{n_i}^{(i)}) && \text{(by Lemma 5.10)} \\ &= \max_{1 \leq i \leq n} (j_{n_i+1}^{(i)}) = j_{k+1} \end{aligned}$$

which concludes the induction. Hence the sequence $(r_k)_{k \in \mathbb{N}}$ reaches 0 in at most t steps, thus in at most than $|L|(|\mathbb{F}_{\mathcal{G}}|^2 + 2)$ steps. \square

Theorem 5.13. *Let \mathcal{G} be an SPTG. For all locations ℓ , $\text{Val}_{\mathcal{G}}(\ell)$ has at most $O((\Pi^{\text{tr}})^4 |L|^9)$ cutpoints.*

Proof. By using the notations of Lemma 5.12, it suffices to show that the number of cutpoints of $\text{Val}_{\mathcal{G}}(\ell)$ in the interval $[r_{i+1}, r_i]$ (with i from 1 to $j-1 \leq |L|(|\mathbb{F}_{\mathcal{G}}|^2 + 2) - 1 = O(\Pi^{\text{tr}^2} |L|^5)$) is at most $O((\Pi^{\text{tr}})^2 |L|^4)$. However, on such an interval, we know that the value function $\text{Val}_{\mathcal{G}}(\ell)$ is equal to $\text{Val}_{\mathcal{G}_{L', r_i}}(\ell)$. But \mathcal{G}_{L', r_i} is a game where all locations are urgent, and thus by Proposition 4.7, its number of cutpoints is indeed bounded by $O((\Pi^{\text{tr}})^2 |L|^4)$. \square

5.4. Algorithms to compute the value function. The finite optimality of SPTGs allows us to compute the value functions. The proof of Theorem 5.9 suggests a *recursive* algorithm to do so: from an SPTG \mathcal{G} with minimal non-urgent location ℓ^* , solve recursively $\mathcal{G}^{\ell^*, 1}$, $\mathcal{G}^{\ell^*, \text{left}_{\ell^*}(1)}$, $\mathcal{G}^{\ell^*, \text{left}_{\ell^*}(\text{left}_{\ell^*}(1))}$, etc. handling the base case where all locations are urgent with Algorithm 1. While our results above show that this is correct and terminates with a pseudo-polynomial time complexity, we propose instead to solve—without the need for recursion—the sequence of games $\mathcal{G}_{L \setminus (L_u \cup L_f), 1}$, $\mathcal{G}_{L \setminus (L_u \cup L_f), \text{left}_{L \setminus (L_u \cup L_f)}(1)}$, \dots i.e. making all

Algorithm 2: solve(\mathcal{G})

```

Input: SPTG  $\mathcal{G} = (L_{\text{Min}}, L_{\text{Max}}, L_f, L_u, \varphi, \Delta, \pi)$ 
1  $\mathbf{f} = (f_\ell)_{\ell \in L} := \text{solveInstant}(\mathcal{G}, 1)$  /*  $f_\ell: \{1\} \rightarrow \overline{\mathbb{R}}$  */
2  $r := 1$ 
3 while  $0 < r$  do /* Invariant:  $f_\ell: [r, 1] \rightarrow \overline{\mathbb{R}}$  */
4    $\mathcal{G}' := \text{wait}(\mathcal{G}, r, \mathbf{f}(r))$  /*  $r$ -SPTG  $\mathcal{G}' = (L_{\text{Min}}, L_{\text{Max}}, L'_f, L'_u, \varphi', T', \pi')$  */
5    $L'_u := L'_u \cup L$  /* every location is made urgent */
6    $b := r$ 
7   repeat /* Invariant:  $f_\ell: [b, 1] \rightarrow \overline{\mathbb{R}}$  */
8      $a := \max(\text{PossCP}_{\mathcal{G}'} \cap [0, b])$ 
9      $\mathbf{x} = (x_\ell)_{\ell \in L} := \text{solveInstant}(\mathcal{G}', a)$  /*  $x_\ell = \text{Val}_{\mathcal{G}'}(\ell, a)$  */
10    if  $\forall \ell \in L_{\text{Min}} \frac{f_\ell(b) - x_\ell}{b - a} \leq -\pi(\ell) \wedge \forall \ell \in L_{\text{Max}} \frac{f_\ell(b) - x_\ell}{b - a} \geq -\pi(\ell)$  then
11      foreach  $\ell \in L$  do  $f_\ell := \left( \nu \in [a, b] \mapsto f_\ell(b) + (\nu - b) \frac{f_\ell(b) - x_\ell}{b - a} \right) \triangleright f_\ell$ 
12       $b := a$ ;  $\text{stop} := \text{false}$ 
13    else  $\text{stop} := \text{true}$ 
14  until  $b = 0$  or  $\text{stop}$ 
15   $r := b$ 
16 return  $\mathbf{f}$ 

```

locations urgent at once. Lemma 5.12 explains why this sequence of games correctly computes the value function of \mathcal{G} and terminates after a pseudo-polynomial number of steps.

Algorithm 2 implements these ideas. Each iteration of the **while** loop computes a new game in the sequence $\mathcal{G}_{L \setminus (L_u \cup L_f), 1}$, $\mathcal{G}_{L \setminus (L_u \cup L_f), \text{left}_{L \setminus (L_u \cup L_f)}(1)}$, \dots ; solves it thanks to `solveInstant`; and thus computes a new portion of $\text{Val}_{\mathcal{G}}$ on an interval on the left of the current point $r \in [0, 1]$. More precisely, the vector $(\text{Val}_{\mathcal{G}}(\ell, 1))_{\ell \in L}$ is first computed in line 1. Then, the algorithm enters the **while** loop, and the game \mathcal{G}' obtained when reaching line 6 is $\mathcal{G}_{L \setminus (L_u \cup L_f), 1}$. Then, the algorithm enters the **repeat** loop to analyse this game. Instead of building the whole value function of \mathcal{G}' , Algorithm 2 builds only the parts of $\text{Val}_{\mathcal{G}'}$ that coincide with $\text{Val}_{\mathcal{G}}$. It proceeds by enumerating the possible cutpoints a of $\text{Val}_{\mathcal{G}'}$, starting in r , by decreasing clock values (line 8), and computes the value of $\text{Val}_{\mathcal{G}'}$ in each cutpoint thanks to `solveInstant` (line 9), which yields a new piece of $\text{Val}_{\mathcal{G}'}$. Then, the **if** in line 10 checks whether this new piece coincides with $\text{Val}_{\mathcal{G}}$, using the condition given by Proposition 5.2. If it is the case, the piece of $\text{Val}_{\mathcal{G}'}$ is added to f_ℓ (line 11); **repeat** is stopped otherwise. When exiting the **repeat** loop, variable b has value $\text{left}_{L \setminus (L_u \cup L_f)}(1)$. Hence, at the next iteration of the **while** loop, $\mathcal{G}' = \mathcal{G}_{L \setminus (L_u \cup L_f), \text{left}_{L \setminus (L_u \cup L_f)}(1)}$ when reaching line 6. By continuing this reasoning inductively, one concludes that the successive iterations of the **while** loop compute the sequence $\mathcal{G}_{L \setminus (L_u \cup L_f), 1}$, $\mathcal{G}_{L \setminus (L_u \cup L_f), \text{left}_{L \setminus (L_u \cup L_f)}(1)}$, \dots as announced, and rebuilds $\text{Val}_{\mathcal{G}}$ from them.

Termination of the **while** loop in pseudo-polynomially many steps is then ensured by Lemma 5.12. Similarly, the termination of the internal **repeat** loop is ensured by the at most pseudo-polynomial number of possible cutpoints and the `stop` variable. As each of the non-trivial calls requires at most pseudo-polynomial time, Algorithm 2 finishes in

pseudo-polynomial time, in total. Note that some SPTGs indeed have a pseudo-polynomial number of cutpoints [FIJS20] (even in the case of only non-negative prices), which shows that our bound is asymptotically tight.

Remark 5.14. The pseudo-polynomial lower-bound on the number of cutpoints shown in [FIJS20] helps getting a PSPACE-hardness of the value problem consisting in deciding whether the value $\text{Val}_{\mathcal{G}}(\ell, 0)$ is below a given rational threshold. We might thus wonder whether our upper-bound techniques help closing the gap. Unfortunately, this does not seem to be the case. Indeed, even if we transform our algorithm to only record the current values $(f_{\ell}(r))_{\ell \in L}$ of the value function, we are not able to obtain that such values can be stored in polynomial space (should we obtain such a result, it would easily imply a polynomial space algorithm to compute the initial values $(f_{\ell}(0))_{\ell \in L}$, since the rest of the algorithm, in particular `solveInstant`, performs in polynomial space). The problem comes from the growth of the various coefficients appearing during the algorithm, in particular the granularity of the rational cutpoints we encounter through the computation. Though unrealistic, if the cutpoint on the left of r was always in the middle of the interval $[0, r]$, cutpoints would have the shape $1/2^x$ with x an integer bounded pseudo-polynomially. Unfortunately, the denominator of this ratio cannot be stored in polynomial space. Thus, getting a polynomial space algorithm to solve SPTGs requires a better understanding of the *granularity* of cutpoints, and not only a bound on their *number*.

Example 5.15. Figure 11 shows the value functions of the SPTG of Figure 1. Here is how Algorithm 2 obtains those functions. During the first iteration of the **while** loop, the algorithm computes the correct value functions until the cutpoint $\frac{3}{4}$: in the *repeat* loop, at first $a = 9/10$ but the slope in ℓ_1 is smaller than the slope that would be granted by waiting, as depicted in Figure 1. Then, $a = 3/4$ where the algorithm gives a slope of value -16 in ℓ_2 while the weight of this location of Max is -14 . During the first iteration of the **while** loop, the inner **repeat** loop thus ends with $r = 3/4$. The next iterations of the **while** loop end with $r = \frac{1}{2}$ (because ℓ_1 does not pass the test in line 10); $r = \frac{1}{4}$ (because of ℓ_2) and finally with $r = 0$, giving us the value functions on the entire interval $[0, 1]$.

6. TOWARDS MORE COMPLEX PTGS

In [BLMR06, Rut11, HIJM13], *general* PTGs with *non-negative weights* are solved by reducing them to a finite sequence of SPTGs, by eliminating guards and resets. It is thus natural to try and adapt these techniques to our general case, in which case Algorithm 2 would allow us to solve *general PTGs with arbitrary weights*. Let us explain where are the difficulties of such a generalisation.

The technique used to remove strict guards from the transitions of the PTGs, i.e. guards of the form $(a, b]$, $[b, a)$ or (a, b) with $a, b \in \mathbb{N}$, consists in enhancing the locations with regions while keeping an equivalent game. This technique *can* be adapted to arbitrary weights. Formally, let $\mathcal{G} = (L_{\text{Min}}, L_{\text{Max}}, L_f, L_u, \varphi, \Delta, \pi)$ be a PTG. We define the region-PTG of \mathcal{G} as $\mathcal{G}' = (L'_{\text{Min}}, L'_{\text{Max}}, L'_f, L'_u, \varphi', \Delta', \pi')$ where:

- $L'_{\text{Min}} = \{(\ell, I) \mid \ell \in L_{\text{Min}}, I \in \text{Reg}_{\mathcal{G}}\}$;
- $L'_{\text{Max}} = \{(\ell, I) \mid \ell \in L_{\text{Max}}, I \in \text{Reg}_{\mathcal{G}}\}$;
- $L'_f = \{(\ell, I) \mid \ell \in L_f, I \in \text{Reg}_{\mathcal{G}}\}$;
- $L'_u = \{(\ell, I) \mid \ell \in L_u, I \in \text{Reg}_{\mathcal{G}}\}$;




Figure 11: Value functions of the SPTG of Figure 1

- for all $(\ell, I) \in L'_f$, if I is a singleton $\{a\}$ then $\varphi'_{\ell, I}(a) = \varphi_\ell(a)$, otherwise I is an interval (a, b) , we then define for $x \in I$, $\varphi'_{\ell, I}(x) = \varphi_\ell(x)$ and extend $\varphi'_{\ell, I}$ on the borders of I by continuity;
- transitions given by

$$\Delta' = \left\{ ((\ell, I), \overline{I_g \cap I}, R, (\ell', I')) \mid (\ell, I_g, R, \ell') \in \Delta, I' = \begin{cases} I & \text{if } R = \perp \\ \{0\} & \text{otherwise} \end{cases} \right\} \cup \text{WaitTr}$$

with

$$\begin{aligned} \text{WaitTr} = & \{ ((\ell, (M_k, M_{k+1})), \{M_{k+1}\}, \perp, (\ell, \{M_{k+1}\})) \mid \ell \in L, (M_k, M_{k+1}) \in \text{Reg}_G \} \\ & \cup \{ ((\ell, \{M_k\}), \{M_k\}, \perp, (\ell, (M_k, M_{k+1}))) \mid \ell \in L, (M_k, M_{k+1}) \in \text{Reg}_G \}; \end{aligned}$$

- $\forall (\ell, I) \in L'$, $\pi'(\ell, I) = \pi(\ell)$; and $\forall \delta' \in \Delta'$, we let $\pi'(\delta')$ being the maximal (resp. minimal) weight of a transition of Δ giving rise to δ' in the definition above, knowing that transitions coming from WaitTr are given weight 0, if ℓ belongs to Max (resp. Min).⁸

⁸Indeed, notice that a transition $((\ell, I), I'', R, (\ell', I')) \in \Delta'$ can be originated from two different transitions (ℓ, I_g, R, ℓ') and (ℓ, I'_g, R, ℓ') of Δ if I_g and I'_g both intersect I .

It is easy to check that the region-PTG fulfils certain *invariants*. In all configurations $((\ell, \{M_k\}), \nu)$ reachable from the clock value 0, the clock value ν is M_k . More interestingly, in all configurations $((\ell, (M_k, M_{k+1})), \nu)$ reachable from the clock value 0, the clock value ν is in $[M_k, M_{k+1}]$, and not only in (M_k, M_{k+1}) as one might expect. Intuitively, we rely on $\nu = M_k$, for example, to denote a configuration of the original game with a clock value arbitrarily close to M_k , but greater than M_k . The game can thus take transitions with guard $x > M_k$, but cannot take transitions with guard $x = M_k$ anymore.

Lemma 6.1. *Let \mathcal{G} be a PTG, and \mathcal{G}' be its region-PTG defined as before. For $(\ell, I) \in L \times \text{Reg}_{\mathcal{G}}$ and $\nu \in I$, $\text{Val}_{\mathcal{G}}(\ell, \nu) = \text{Val}_{\mathcal{G}'}((\ell, I), \nu)$. Moreover, we can transform an ε -optimal strategy of \mathcal{G}' into an ε' -optimal strategy of \mathcal{G} with $\varepsilon' < 2\varepsilon$ and vice-versa.*

Proof. Intuitively, the proof consists in replacing strategies of \mathcal{G}' where players can play on the borders of regions, by strategies of \mathcal{G} that play increasingly close to the border as time passes. If played close enough, the loss created can be chosen as small as we want.

Formally, let \mathcal{G} be a PTG, \mathcal{G}' be its region-PTG. First, for $\varepsilon > 0$, we create a transformation g of the plays of \mathcal{G}' which do not end with a waiting transition to the plays of \mathcal{G} . It is defined by induction on the length n of the plays so that for a play ρ of length n we have

- $|\text{Cost}(\rho) - \text{Cost}(g(\rho))| \leq 2\Pi^{\text{loc}}(1 - \frac{1}{2^n})\varepsilon$; and
- there exists $\ell \in L$ and $I \in \text{Reg}_{\mathcal{G}}$ such that $g(\rho)$ and ρ end in the respective locations ℓ and (ℓ, I) , and their clock values are both in I and differ of at most $\frac{1}{2^{n+1}}\varepsilon$.

If $n = 0$, let $\rho = ((\ell, I), \nu)$ be a play of \mathcal{G}' of length 0, then $g(\rho) = (\ell, \nu')$, where $\nu' = \nu \pm \frac{\varepsilon}{2}$ if I is not a singleton and ν is an endpoint of I , and $\nu' = \nu$ otherwise (so that $\nu' \in I$ in every case).

For $n > 0$, we suppose g defined on every play of length at most n which does not end with a waiting transition. Let $\rho = ((q_1, I_1), \nu_1) \xrightarrow{t_1, \delta_1, c_1} \dots \xrightarrow{t_n, \delta_n, c_n} ((q_n, I_n), \nu_n) \xrightarrow{t_{n+1}, \delta_{n+1}, c_{n+1}} ((q_{n+1}, I_{n+1}), \nu_{n+1})$ with $\delta_{n+1} \notin \text{WaitTr}$. Let $last = \max(\{k \leq n \mid tr_k \notin \text{WaitTr}\})$ (with $\max \emptyset = 0$). Then, by induction, there exists $\rho' = (q_1, \nu_1) \rightarrow \dots \rightarrow (q_{last+1}, \nu'_{last+1})$ such that

- $g(\rho|_{last}) = \rho'$ (where $\rho|_{last}$ is the prefix of length $last$ of ρ),
- $|\text{Cost}(\rho|_{last}) - \text{Cost}(g(\rho|_{last}))| \leq 2\Pi^{\text{loc}}(1 - \frac{1}{2^{last}})\varepsilon$, and
- $|\nu'_{last+1} - \nu_{last+1}| \leq \frac{1}{2^{last+1}}\varepsilon$.

Then we choose $g(\rho) = \rho' \xrightarrow{t, \delta'_{n+1}, c} (q_{n+1}, \nu'_{n+1})$, with δ'_{n+1} the transition giving rise to δ_{n+1} (with the correct price) in the definition of the region-PTG, and where

- if δ_{n+1} is enabled in configuration $(q_{last+1} = q_n, \nu_n + t_{n+1})$ of \mathcal{G} , then, $t = \nu_n + t_{n+1} - \nu'_{last}$;
- otherwise, as the guards of \mathcal{G}' are contained in the closure¹⁰ of the guards of \mathcal{G} , then there exists $z \in \{1, -1\}$ such that for $t = \nu_n + t_{n+1} - \nu'_{last} + \frac{z\varepsilon}{2^{n+2}}$, δ_{n+1} is enabled in \mathcal{G} and $\nu'_{last} + t$ and $\nu_n + t_{n+1}$ belong to the same region.

⁹Recall that a strategy σ_{Min} of Min is ε -optimal from location ℓ in \mathcal{G} if $\text{Price}(\ell, \sigma_{\text{Min}}) \leq \text{Val}_{\mathcal{G}}(\ell) + \varepsilon$.

¹⁰By closure, we mean that, for example, a guard of the form $x > 1$ becomes $x \geq 1$ in the region $[1, 2]$.

Thus, in both cases, $|\nu_{n+1} - \nu'_{n+1}| \leq \frac{\varepsilon}{2^{n+2}}$ and $\nu_{n+1} \neq \nu'_{n+1}$ iff I is not a singleton, ν_{n+1} is on a border, ν'_{n+1} is close to this border and δ_{n+1} does not contain a reset. Moreover,

$$\begin{aligned}
|\text{Cost}(\rho) - \text{Cost}(g(\rho))| &= |\text{Cost}(\rho|_{last}) + (\nu_{n+1} - \nu_{last})\pi(q_{last}) + \pi(\delta_{n+1}) - \text{Cost}(g(\rho))| \\
&\leq |\text{Cost}(\rho|_{last}) - \text{Cost}(g(\rho|_{last}))| \\
&\quad + |(\nu_{n+1} - \nu_{last})\pi(q_{last}) + \pi(\delta_{n+1}) + \text{Cost}(g(\rho|_{last})) - \text{Cost}(g(\rho))| \\
&\leq 2\Pi^{\text{loc}}(1 - \frac{1}{2^{last}})\varepsilon + |(\nu'_{last} - \nu_{last})\pi(q_{last}) + (\nu_{n+1} - \nu'_{n+1})\pi(q_{last})| \\
&\leq 2\Pi^{\text{loc}}(1 - \frac{1}{2^{last}})\varepsilon + \left| \frac{\varepsilon}{2^{last+1}}\pi(q_{last}) \right| + \left| \frac{\varepsilon}{2^{n+2}}\pi(q_{last}) \right| \\
&\leq 2\Pi^{\text{loc}}(1 - \frac{1}{2^{last}})\varepsilon + \frac{\Pi^{\text{loc}}\varepsilon}{2^{last+1}} + \frac{\Pi^{\text{loc}}\varepsilon}{2^{n+2}} \\
&\leq 2\Pi^{\text{loc}}(1 - \frac{1}{2^{last+1}})\varepsilon \\
&\leq 2\Pi^{\text{loc}}(1 - \frac{1}{2^{n+1}})\varepsilon.
\end{aligned}$$

Let σ_{Min} be a strategy of Min in \mathcal{G} . Using the transformation g , we will build by induction a strategy σ'_{Min} in \mathcal{G}' such that, for all plays ρ whose last transition does not belong to WaitTr and conforming with σ'_{Min} : $g(\rho)$ conforms with σ_{Min} .

Let ρ be a play of \mathcal{G}' whose last transition does not belong to WaitTr such that $g(\rho)$ conforms with σ_{Min} (which is the case of all plays of length 0). Then, ρ and $g(\rho)$ end in locations (q, I) and q respectively.

- If ρ ends in a configuration of Max , then the choice of the next (t, δ) -transition does not depend on σ_{Min} or σ'_{Min} . Let (t, δ) be a choice of Max in \mathcal{G}' with cost c . If δ belongs to WaitTr , then the new configuration also belongs to Max where it will make another choice. Let ρ' be the extension of ρ until the first transition δ' such that $\delta' \notin \text{WaitTr}$. The play $g(\rho')$ conforms with σ_{Min} as the configuration where $g(\rho)$ ends is controlled by Max and $g(\rho')$ only has one more transition than $g(\rho)$.
- If ρ ends in a configuration of Min , then there exists t, δ, c, q', ν' such that $g(\rho) \xrightarrow{t, \delta, c} (q', \nu')$ conforms with σ_{Min} . As taking a waiting transition does not change the ownership of the configuration, we consider here multiple successive choices of Min as one choice: $\sigma'_{\text{Min}}(\rho)$ is such that $\rho' = \rho \xrightarrow{t_1, \delta_1, c_1} \dots \xrightarrow{t_k, \delta_k, c_k} ((q, I''), \nu) \xrightarrow{t_{k+1}, \delta, c_{k+1}} ((q', I'), \nu')$ where $\forall i \leq k, \delta_i \in \text{WaitTr}$ conforms with σ'_{Min} . This is possible as if δ is allowed in a configuration (q, ν) in \mathcal{G} then it is allowed too in a configuration $((q, I), \nu)$ with the appropriate I . Then $g(\rho') = g(\rho) \xrightarrow{\delta, tr, c} (q', \nu')$, thus $g(\rho')$ conforms with σ_{Min} .

As no completed plays of \mathcal{G}' end with a transition of WaitTr , every completed play ρ conforming with σ'_{Min} verifies that $g(\rho)$ is a completed play conforming with σ_{Min} . Moreover, the time valuation of ρ and $g(\rho)$ differ by at most ε and belong to the same interval (potentially on one border), thus by definition of φ' the difference in final cost is bounded by $K_{fin}\varepsilon$ where K_{fin} is the greatest absolute value of the slopes appearing in the piecewise affine functions within φ . Thus for every configuration s , $\text{Price}_{\mathcal{G}'}(s, \sigma'_{\text{Min}}) \leq \text{Price}_{\mathcal{G}}(s, \sigma_{\text{Min}}) + (2\Pi^{\text{loc}} + K_{fin})\varepsilon$. Therefore $\text{Val}_{\mathcal{G}'}(s) \leq \text{Val}_{\mathcal{G}}(s)$.

Reciprocally, let σ'_{Min} be a strategy of Min in \mathcal{G}' . We will now build by induction a strategy σ_{Min} in \mathcal{G} such that for all plays ρ conforming with σ_{Min} , there exists a play in $g^{-1}(\rho)$ that conforms with σ'_{Min} .

Let ρ be a play of \mathcal{G} conforming with σ_{Min} such that there exists $\rho' \in g^{-1}(\rho)$ conforming with σ'_{Min} (which is the case of all plays of length 0). Plays ρ' and ρ end in the configurations $((q, I), \nu')$ and (q, ν) respectively.

- If ρ ends in configuration of Max, then the choice does not depend on σ_{Min} or σ'_{Min} . Let (t, δ) be a choice of Max in \mathcal{G} with cost c and let $\tilde{\rho}$ be the extension of ρ by this choice. There exists $(t_1, \delta_1, c_1), \dots, (t_{k+1}, \delta_{k+1}, c_{k+1})$ such that $\forall i \leq k, \delta_i \in \text{WaitTr}, \delta_{k+1} = \delta$ and $\sum_{i=1}^{k+1} t_i = \nu + t - \nu'$. Let $\rho_c = \rho' \xrightarrow{t_1, \delta_1, c_1} \dots \xrightarrow{t_k, \delta_k, c_k} ((q, I''), \nu_k) \xrightarrow{t_{k+1}, \delta_{k+1}, c_{k+1}} ((q', I'), \nu_{k+1})$, then ρ_c conforms with σ'_{Min} (as Min did not take a single decision) and $g(\rho_c) = \tilde{\rho}$.
- If ρ ends in a configuration of Min, then there exists a play $\rho_c = \rho \xrightarrow{t_1, \delta_1, c_1} \dots \xrightarrow{t_k, \delta_k, c_k} ((q, I''), \nu_k) \xrightarrow{t_{k+1}, \delta_{k+1}, c_{k+1}} ((q', I'), \nu_{k+1})$ such that ρ_c conforms with σ'_{Min} . We choose $\sigma_{\text{Min}}(\rho) = (t, \delta)$ such that for the adequate cost c , $g(\rho_c) = \rho \xrightarrow{t, \delta, c} (q', \nu')$. This is possible as $t + \nu' \in I''$.

Every completed play ρ conforming with σ_{Min} verifies $\exists \rho' \in g^{-1}(\rho)$ conforming with σ_{Min} . Thus, taking the final cost function into account as before, for every configuration s , $\text{Price}_{\mathcal{G}}(s, \sigma_{\text{Min}}) \leq \text{Price}_{\mathcal{G}'}(s, \sigma_{\text{Min}}) + (2\Pi^{\text{loc}} + K_{\text{fin}})\varepsilon$. Therefore $\text{Val}_{\mathcal{G}'}(s) \geq \text{Val}_{\mathcal{G}}(s)$. Hence $\text{Val}_{\mathcal{G}'}(s) = \text{Val}_{\mathcal{G}}(s)$. \square

The technique used in [BLMR06, Rut11, HIJM13] to remove resets from PTGs, however, consists in *bounding* the number of clock resets that can occur in each play following an optimal strategy of Min or Max. Then, the PTG can be *unfolded* into a *reset-acyclic* PTG with the same value. By reset-acyclic, we mean that no cycles in the configuration graph visit a transition with a reset. This reset-acyclic PTG can be decomposed into a finite number of components that contain no reset and are linked by transitions with resets. These components can be solved iteratively, from the bottom to the top, turning them into SPTGs. Thus, if we *assume* that the PTGs we are given as input *are* reset-acyclic, we can solve them in *pseudo-polynomial time*, and show that their value functions are cost functions with at most a pseudo-polynomial number of cutpoints, using our techniques.

In [BLMR06] the authors showed that with one-clock PTG and non-negative weights only we could bound the number of resets by the number of locations, without changing the value functions. Unfortunately, these arguments do not hold for arbitrary weights, as shown by the PTG in Figure 12. In that PTG, we claim that $\text{Val}(\ell_0) = 0$; that Min has no optimal strategies, but a family of ε -optimal strategies $\sigma_{\text{Min}}^\varepsilon$ each with value ε ; and that each $\sigma_{\text{Min}}^\varepsilon$ requires *memory whose size depends on ε* and might *yield a play visiting at least $1/\varepsilon$ times the reset* between ℓ_1 and ℓ_0 (hence the number of resets cannot be bounded). For all $\varepsilon > 0$, $\sigma_{\text{Min}}^\varepsilon$ consists in: waiting $1 - \varepsilon$ time units in ℓ_0 , then going to ℓ_1 during the $\lceil 1/\varepsilon \rceil$ first visits to ℓ_0 ; and to go directly to ℓ_f afterwards. Against $\sigma_{\text{Min}}^\varepsilon$, Max has several possible choices:

- (1) either wait $\eta \in [0, \varepsilon]$ time units in ℓ_1 , wait $\varepsilon - \eta$ time units in ℓ_2 , then reach ℓ_f ; or
- (2) wait ε time unit in ℓ_1 to have the clock equal to 1, and force the cycle by going back to ℓ_0 , where the game will wait for Min's next move.

Thus, all plays according to $\sigma_{\text{Min}}^\varepsilon$ will visit a sequence of locations which is either of the form $\ell_0(\ell_1\ell_0)^k\ell_1\ell_2\ell_f$, with $0 \leq k < \lceil 1/\varepsilon \rceil$; or of the form $\ell_0(\ell_1\ell_0)^{\lceil 1/\varepsilon \rceil}\ell_f$. In the former case, the price of the play will be $-k\varepsilon - \eta + (\varepsilon - \eta) = -(k-1)\varepsilon - 2\eta \leq \varepsilon$; in the latter, $-\varepsilon(\lceil 1/\varepsilon \rceil) + 1 \leq 0$. This shows that $\text{Val}(\ell_0) = 0$, but there are no optimal strategies as none




Figure 12: A PTG where the number of resets in optimal plays cannot be bounded a priori.

of these strategies allow one to guarantee a price of 0 (neither does the strategy that waits 1 time unit in ℓ_0).

If bounding the number of resets is not possible in the general case, it could be done if one adds constraints on the cycles of the game. This kind of restriction was used in [BCR14] where the authors introduce the notion of *robust* games (and a more restrictive one of *divergent* games was used in [BGMR17]). Such games require among other things that there exists $\kappa > 0$ such that every play starting and ending in the same pair location and time region has either a positive cost or a cost smaller than $-\kappa$. Here we require a less powerful assumption as we put this restriction only on cycles containing a reset.

Definition 6.2. Given $\kappa > 0$, a κ -negative-reset-acyclic PTG (κ -NRAPTG) is a PTG where for every location $\ell \in L$ and every cyclic finite play ρ starting and ending in $(\ell, 0)$, either $\text{Cost}(\rho) \geq 0$ or $\text{Cost}(\rho) < -\kappa$.

The PTG of Figure 12 is not a κ -NRAPTG for any $\kappa > 0$ as the play $(\ell_0, 0) \xrightarrow{0} (\ell_1, 1 - \kappa/2) \xrightarrow{-\kappa/2} (\ell_0, 0)$ is a cycle containing a reset and with a negative cost strictly greater than $-\kappa$. On the contrary, in Figure 13 we show a 1-NRAPTG and its region PTG. Here, every cycle containing a reset is between ℓ_0 and ℓ_1 and such cycles have at most cost -1 . The value of this PTG is 0 but no strategies for Max can achieve it because of the guard $x > 0$ on the transition from ℓ_1 to ℓ_f . As this guard is not strict anymore in the region PTG, both players have an optimal strategy in this game (this is not always the case).

In order to bound the number of resets of a κ -NRAPTG, we first prove a bound on the value of such games, that will be useful in the following. We let $k = |\text{Reg}_{\mathcal{G}}|$ be the number of regions. Recall that M is a bound on the valuations taken by the clock in \mathcal{G} , as discussed on page 8.

Lemma 6.3. For all κ -NRAPTGs \mathcal{G} and $(\ell, \nu) \in \text{Conf}_{\mathcal{G}}$: either $\text{Val}_{\mathcal{G}}(\ell, \nu) \in \{-\infty, +\infty\}$, or

$$-|L|M\Pi^{\text{loc}} - |L|^2(|L| + 2)\Pi^{\text{tr}} - \Pi^{\text{fin}} \leq \text{Val}_{\mathcal{G}}(\ell, \nu) \leq |L|M\Pi^{\text{loc}} + |L|k\Pi^{\text{tr}} + \Pi^{\text{fin}}.$$

Proof. Consider the case where $\text{Val}_{\mathcal{G}}(\ell, \nu) \notin \{-\infty, +\infty\}$. Let $\kappa > 2\varepsilon > 0$. Then, there exist σ_{Min} and σ_{Max} ε -optimal strategies for Min and Max, respectively.

Let $\sigma_{\text{Min}}^{\neg c}$ be any memoryless strategy of Min in the reachability timed game induced by \mathcal{G} such that no play consistent with $\sigma_{\text{Min}}^{\neg c}$ goes twice in the same couple (location, region). If such




Figure 13: A 1-NRAPTG and its region PTG (some guards removed for better readability)

a strategy does not exist, as the clock constraints are the same during the first and second occurrences of this couple, Max can enforce the cycle infinitely often, thus the reachability game is winning for Max and the value of \mathcal{G} is $+\infty$. Let us note $\rho = \text{CPlay}((\ell, \nu), \sigma_{\text{Min}}^{-c}, \sigma_{\text{Max}})$. By ε -optimality of σ_{Max} , $\text{Price}(\rho) \geq \text{Val}_{\mathcal{G}}(\ell, \nu) - \varepsilon$. Let $\text{Cost}^{\text{tr}}(\rho)$ be the price of ρ due to the weights of the transitions, and $\text{Cost}^{\text{loc}}(\rho)$ be the weight due to the time elapsed in the locations of the game: $\text{Cost}(\rho) = \text{Cost}^{\text{tr}}(\rho) + \text{Cost}^{\text{loc}}(\rho)$. As there are no cycles in the game according to couples (location, region), there are at most $|L|k$ transitions, thus $\text{Cost}^{\text{tr}}(\rho) \leq |L|k\Pi^{\text{tr}}$. Moreover, the absence of cycles also implies that we do not take two transitions with a reset ending in the same location or one transition with a reset ending in the initial location, thus we take at most $|L| - 1$ such transitions. Therefore at most $|L|M$ units of time elapsed and $\text{Cost}^{\text{loc}}(\rho) \leq |L|M\Pi^{\text{loc}}$. Adding the final cost, this implies that

$$\text{Val}_{\mathcal{G}}(\ell, \nu) - \varepsilon \leq \text{Price}(\rho) \leq |L|M\Pi^{\text{loc}} + |L|k\Pi^{\text{tr}} + \Pi^{\text{fin}}.$$

By taking the limit of ε towards 0, we obtain the announced upper bound.

We now prove the lower bound on the value. To that extent, consider now the completed play $\rho = \text{CPlay}((\ell, \nu), \sigma_{\text{Min}}, \sigma_{\text{Max}})$. We have that $\text{Price}(\rho) \leq \text{Val}_{\mathcal{G}}(\ell, \nu) + \varepsilon$.

We want to lower bound the price of ρ , therefore non-negative cycles can be safely ignored. Let us show that there are no negative cycles around a transition with a reset. If it was the case, since the game is a κ -NRAPTG, this cycle has cost at most $-\kappa$. Since the strategy σ_{Max} is ε -optimal, and $\kappa > \varepsilon$, it is not possible that σ_{Max} decides alone to take this bad cycle. Therefore, σ_{Min} has the capability to enforce this cycle, and to exit it (otherwise, Max would keep it inside to get value $+\infty$): but then, Min could decide to cycle as long as it wants, then guaranteeing a value as low as possible, which contradicts the fact that $\text{Val}(\ell, \nu) \notin \{-\infty, +\infty\}$. Therefore, the only cycles in ρ around transitions with resets, are non-negative cycles. This implies that its price is bounded below by the price of a sub-play obtained by removing the cycles in ρ .

We now consider a play where each reset transition is taken at most once in ρ , and lower-bound its price.

If ρ contains a cycle around a location $\ell' \in L_{\text{Max}}$ without reset transitions, this cycle has the form $(\ell', \nu') \xrightarrow{c'} (\ell'', \nu' + t) \cdots \xrightarrow{c''} (\ell', \nu'')$ with $\nu'' \geq \nu'$, followed in ρ by a transition towards configuration $(\ell''', \nu'' + t')$. Thus, another strategy for Max could have consisted in skipping the cycle by choosing as delay in the first location ℓ' , $\nu'' - \nu' + t'$ instead of t . This would get a new strategy that cannot make the price increase above $\text{Val}_{\mathcal{G}}(\ell, \nu) + \varepsilon$, since it is still playing against an ε -optimal strategy of Min. Therefore, we can consider the sub-play ρ_f of ρ where all such cycles are removed: we still have $\text{Price}(\rho_f) \leq \text{Val}_{\mathcal{G}}(\ell, \nu) + \varepsilon$.

Suppose now that ρ_f contains a cycle around a location $\ell' \in L_{\text{Min}}$ without reset transitions, of the form $(\ell', \nu') \xrightarrow{c'} (\ell'', \nu' + t) \cdots \xrightarrow{c''} (\ell', \nu'')$ with ν' and ν'' in the same region, composed of Min's locations only, and followed in ρ by a transition towards configuration $(\ell''', \nu'' + t')$. Then, the transition weight of this cycle is non-negative, otherwise Min could enforce this cycle it entirely controls, while letting only a bounded time pass (smaller and smaller as the number of cycles grow). This is not possible.

Therefore, we have that two occurrences of a same Max's location in ρ_f are separated by a reset transition and two occurrences of a same Min's couple (location, region) are either separated by a reset or by a Max's location. As there are at most $|L| - 1$ resets, $|L|$ locations of Max and $|L|k$ couples (location, region) for Min, ρ_t contains at most $|L|^2$ locations of Max and $|L|k(|L|^2 + |L| - 1 + 1)$ locations of Min, which makes for at most $|L|^2(|L|k + k + 1)$ locations. Thus $\text{Cost}^{\text{loc}}(\rho_t) \geq -|L|^2(|L|k + k + 1)\Pi^{\text{loc}}$. Moreover, as at most $|L| - 1$ resets are taken in ρ_f and that the game is bounded by M , $\text{Cost}^{\text{loc}}(\rho_f) \geq -|L|M\Pi^{\text{loc}}$. Adding the final cost, this implies that

$$\text{Val}_{\mathcal{G}}(\ell, \nu) + \varepsilon \geq \text{Cost}^{\text{loc}}(\rho_f) + \text{Cost}^{\text{tr}}(\rho_t) \geq -|L|M\Pi^{\text{loc}} - |L|^2(|L|k + k + 1)\Pi^{\text{tr}} - \Pi^{\text{fin}}.$$

Taking the limit when ε tends to 0, we obtain the desired lower bound. \square

Using this bound on the value of a κ -NRAPTG, one can give a bound on the number of cycles needed to be allowed. The idea is that if a reset is taken twice and the generated cycle has positive cost, either Min can modify its strategy so that it does not take this cycle or the value of the game is $+\infty$ as Max can prevent Min from reaching a final location. On the contrary if the cycle has negative cost, then by definition of a κ -NRAPTG, this cost is less than $-\kappa$. Thus by allowing enough such cycles, as we have bounds on the values of the game, we know when we will have enough cycles to get under the lower bound of the value of the game. By solving the copies of the game, if we reach a value that is smaller than the lower bound of the value, then it means that the value is $-\infty$.

Lemma 6.4. *For all $\kappa > 0$, the value of a κ -NRAPTG can be computed by solving $\lceil 2n(\text{Val}^{\text{sup}} - \text{Val}_{\text{inf}})/\kappa \rceil$ PTGs without resets and using the same set of guards, where Val^{sup} and Val_{inf} are the upper and lower bounds of the value of the game given by Lemma 6.3. Moreover, from ε -optimal strategies on those k games, we can build $\varepsilon \lceil 2n(\text{Val}^{\text{sup}} - \text{Val}_{\text{inf}})/\kappa \rceil$ -optimal strategies in the original game.*

Solving the PTGs without resets can be done by using the same algorithms as the one described before for SPTGs: indeed, since we play in a region-PTG, we can focus on the resolution of a subgame staying in the same region until the final transition, and such a game can even be decomposed into simpler games where every region has length 1, which can then be interpreted as a SPTG. Another possibility would be to rescale the time and the weights in order to transform a region (a, b) into $(a, a + 1)$, avoiding to split the region into $b - a$ different subregions.

The values Val^{sup} and Val_{inf} , and therefore also $\lceil 2n(\text{Val}^{sup} - \text{Val}_{inf})/\kappa \rceil$, are pseudo-polynomial in the size of the original game, which allows us to conclude:

Theorem 6.5. *Let $\kappa > 0$ and \mathcal{G} be a κ -NRAPTG. Then for every location $q \in Q$, the function $\nu \mapsto \text{Val}_{\mathcal{G}}(q, \nu)$ is computable in pseudo-polynomial time and is piecewise-affine with at most a pseudo-polynomial number of cutpoints. Moreover, for every $\varepsilon > 0$, there exist (and we can effectively compute) ε -optimal strategies for both players.*

The robust games defined in [BCR14] restricted to one-clock are a subset of the NRAPTG, therefore their value is computable with the same complexity. While we cannot extend the computation of the value to all (one-clock) PTGs, we can still obtain information on the nature of the value function:

Theorem 6.6. *The value functions of all one-clock PTGs are cost functions with at most a pseudo-polynomial number of cutpoints.*

Proof. Let \mathcal{G} be a one-clock PTG. Let us replace all transitions (ℓ, g, \top, ℓ') resetting the clock by (ℓ, g, \perp, ℓ'') , where ℓ'' is a new final location with $\varphi_{\ell''} = \text{Val}_{\mathcal{G}}(\ell', 0)$ —observe that $\text{Val}_{\mathcal{G}}(\ell, 0)$ exists even if we cannot compute it, so this transformation is well-defined. This yields a reset-acyclic PTG \mathcal{G}' such that $\text{Val}_{\mathcal{G}'} = \text{Val}_{\mathcal{G}}$. The pseudo-polynomial number of cutpoints of reset-acyclic PTG, as for SPTGs, does not depend on the size of final prices (but only on the price of transitions, and the number of locations), which allows us to conclude. \square

As a consequence, in the particular case of non-negative prices only where transitions with a reset can be unfolded to remove cycles in which they are contained, this ensures that the exponential-time algorithms of [BLMR06, Rut11, HIJM13] indeed have a pseudo-polynomial time complexity.

Corollary 6.7. *The value functions of all one-clock PTGs with only non-negative prices can be computed in pseudo-polynomial time.*

7. CONCLUSION

In this work, we study, for the first time, priced timed games with arbitrary weights and one clock, showing how to compute optimal values and strategies in pseudo-polynomial time for the special case of simple games. This complexity result is better than previously obtained results in the case of non-negative weights only [HIJM13, Rut11] (where an exponential complexity was obtained), and we follow different paths to prove termination and (partial) correctness (due to the presence of negative weights). In order to push our algorithm as far as we can, we introduce the class of negative-reset-acyclic games for which we obtain the same result: as a particular case, we can solve all priced timed games with one clock for which the clock is reset in every cycle of the underlying region automaton. As future works, it is appealing to solve the full class of priced timed games with arbitrary weights and one clock. We have shown why our technique seems to break in this more general setting, thus it could be interesting to study the difficult negative cycles without reset as their own, with different techniques.

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. doi:10.1007/978-3-540-27836-8_13.
- [AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994. doi:10.1016/0304-3975(94)90010-8.
- [ALTP04] Rajeev Alur, Salvatore La Torre, and George J. Pappas. Optimal paths in weighted timed automata. *Theoretical Computer Science*, 318(3):297–322, 2004. doi:10.1016/j.tcs.2003.10.038.
- [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. doi:10.1007/s10703-007-0035-4.
- [BBC06] Patricia Bouyer, Thomas Brihaye, and Fabrice Chevalier. Control in o-minimal hybrid systems. In *Proceedings of the Twenty-first Annual IEEE Symposium on Logic In Computer Science (LICS'06)*, pages 367–378. IEEE Computer Society Press, 2006. doi:10.1109/LICS.2006.22.
- [BBJ⁺08] Patricia Bouyer, Thomas Brihaye, Marcin Jurdziński, Ranko Lazić, and Michał Rutkowski. Average-price and reachability-price games on hybrid automata with strong resets. In *Proceedings of the 6th international conference on Formal Modeling and Analysis of Timed Systems (FORMATS'08)*, volume 5215 of *Lecture Notes in Computer Science*, pages 63–77. Springer, 2008. doi:10.1007/978-3-540-85778-5_6.
- [BBM06] Patricia Bouyer, Thomas Brihaye, and Nicolas Markey. Improved undecidability results on weighted timed automata. *Information Processing Letters*, 98(5):188–194, 2006. doi:10.1016/j.ipl.2006.01.012.
- [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 *Lecture Notes in Computer Science*, pages 49–64. Springer, 2005. doi:10.1007/11603009_5.
- [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 *Lecture Notes in Computer Science*, pages 148–160. Springer, 2004. doi:10.1007/978-3-540-30538-5_13.
- [BCJ09] J. Berendsen, T. Chen, and D. Jansen. Undecidability of cost-bounded reachability in priced probabilistic timed automata. In *Theory and Applications of Models of Computation*, volume 5532 of *Lecture Notes in Computer Science*, pages 128–137. Springer, 2009. doi:10.1007/978-3-642-02017-9_16.
- [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, Berlin, Germany, April 15-17*, pages 283–292. Association for Computing Machinery, 2014. doi:10.1145/2562059.2562116.
- [BDG⁺16] Thomas Brihaye, Amit Kumar Dhar, Gilles Geeraerts, Axel Haddad, and Benjamin Monmege. Efficient energy distribution in a smart grid using multi-player games. In Thomas Brihaye, Benoit Delahaye, Nicolas Markey, and Jiří Srba, editors, *Proceedings of the Cassting Workshop on Games for the Synthesis of Complex Systems (Cassting'16) and the 3rd International Workshop on Synthesis of Complex Parameters (SynCoP'16)*, volume 220, pages 1–12, Eindhoven, Netherlands, April 2016. EPTCS. doi:10.4204/EPTCS.220.1.
- [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. doi:10.1007/3-540-45351-2_15.
- [BGH⁺15] Thomas Brihaye, Gilles Geeraerts, Axel Haddad, Engel Lefauchaux, and Benjamin Monmege. Simple priced timed games are not that simple. In Prahladh Harsha and G. Ramalingam, editors, *Proceedings of the 35th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'15)*, volume 45 of *Leibniz International Proceedings in*

- Informatics (LIPICs)*, pages 278–292. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, December 2015. doi:10.4230/LIPICs.FSTTCS.2015.278.
- [BGHM15] Thomas Brihaye, Gilles Geeraerts, Axel Haddad, and Benjamin Monmege. To reach or not to reach? Efficient algorithms for total-payoff games. In Luca Aceto and David de Frutos Escrig, editors, *Proceedings of the 26th International Conference on Concurrency Theory (CONCUR'15)*, volume 42 of *LIPICs*, pages 297–310. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, September 2015. doi:10.4230/LIPICs.CONCUR.2015.297.
- [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.
- [BGK⁺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'13)*, volume 8704 of *Lecture Notes in Computer Science*, pages 560–575. Springer, 2014. doi:10.1007/978-3-662-44584-6_38.
- [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.
- [BJM14] Patricia Bouyer, Samy Jaziri, and Nicolas Markey. On the value problem in weighted timed games. Research Report LSV-14-12, Laboratoire Spécification et Vérification, ENS Cachan, France, October 2014. 24 pages. URL: http://www.lsv.ens-cachan.fr/Publis/RAPPORTS_LSV/PDF/rr-lsv-2014-12.pdf.
- [BLMR06] Patricia Bouyer, Kim G. Larsen, Nicolas Markey, and Jacob Illum Rasmussen. Almost optimal strategies in one-clock priced timed games. In *Proceedings of the 26th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'06)*, volume 4337 of *Lecture Notes in Computer Science*, pages 345–356. Springer, 2006. doi:10.1007/11944836_32.
- [Bou15] Patricia Bouyer. On the optimal reachability problem in weighted timed automata and games. In *Proceedings of the 7th Workshop on Non-Classical Models of Automata and Applications (NCMA'15)*, volume 318 of *books@ocg.at*, pages 11–36. Austrian Computer Society, 2015.
- [dAHM01] Luca de Alfaro, Thomas A. Henzinger, and Rupak Majumdar. Symbolic algorithms for infinite-state games. In *Proceedings of the 12th International Conference on Concurrency Theory (CONCUR'01)*, volume 2154 of *Lecture Notes in Computer Science*, pages 536–550. Springer, 2001. doi:10.1007/3-540-44685-0_36.
- [FIJS20] John Fearnley, Rasmus Ibsen-Jensen, and Rahul Savani. One-clock priced timed games are PSPACE-hard. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Sciences (LICS'20)*, pages 397–409. ACM, 2020. doi:10.1145/3373718.3394772.
- [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. doi:10.1007/978-3-642-40184-8_37.
- [Mar75] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975. doi:10.2307/1971035.
- [MPS95] Oded Maler, Amir Pnueli, and Joseph Sifakis. On the synthesis of discrete controllers for timed systems. In *Proceedings of the 12th Annual Symposium on Theoretical Aspects of Computer Science (STACS'95)*, volume 900 of *Lecture Notes in Computer Science*, pages 229–242. Springer, 1995. doi:10.1007/3-540-59042-0_76.
- [Rut11] Michał Rutkowski. Two-player reachability-price games on single-clock timed automata. In *Proceedings of the 9th Workshop on Quantitative Aspects of Programming Languages (QAPL'11)*, volume 57 of *Electronic Proceedings in Theoretical Computer Science*, pages 31–46, 2011.
- [WT97] Howard Wong-Toi. The synthesis of controllers for linear hybrid automata. In *Proceedings of the 36th IEEE Conference on Decision and Control (CDC'97)*, pages 4607–4612. IEEE Computer Society Press, 1997. doi:10.1109/CDC.1997.649708.