One-Clock Priced Timed Games with Arbitrary Weights⋆

Thomas Brihaye[a], Gilles Geeraerts[b], Axel Haddad[a], Engel Lefaucheux[c], Benjamin Monmege[d,1]

[a] *Université de Mons, Belgium*
[b] *Université libre de Bruxelles, Belgium*
[c] *LSV, ENS Paris-Saclay, Inria, France*
[d] *Aix Marseille Univ, Université de Toulon, LIS, CNRS, Marseille, France*

Abstract

Priced timed games are two-player zero-sum games played on priced timed automata (whose locations and transitions are labeled by weights modeling the price of spending time in a state and executing an action, respectively). The goals of the players are to minimise and maximise the price 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 theirs (the so-called *simple* priced timed games), one can compute, in exponential 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.

Keywords: Priced timed games, Real-time systems, Game theory

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

⋆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)

Email addresses: `thomas.brihaye@umons.ac.be` (Thomas Brihaye), `gigeerae@ulb.ac.be` (Gilles Geeraerts), `axel.haddad@umons.ac.be` (Axel Haddad), `engel.lefaucheux@irisa.fr` (Engel Lefaucheux), `benjamin.monmege@univ-amu.fr` (Benjamin Monmege)

[1]During part of the preparation of this article, the last author was (partially) funded by the DeLTA project (ANR-16-CE40-0007), the SensAS project (INS2I JCJC'17) and the TickTac project (ANR-18).

system, and of the possible actions of a controller, to compute a controller 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 price-rates and transitions are labeled by discrete prices, as in priced timed automata [BFH+01, ALTP04, BBBR07]. These prices allow one to associate a price with each run (or play), which depends on the sequence of transitions traversed by the run, 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: he tries to avoid the target locations, and, if not possible, to maximise the accumulated price up to the first visit of a target location. To reflect these objectives, we define the upper value $\overline{\mathsf{Val}}$ 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 $\underline{\mathsf{Val}}$ returns the greatest price that Max can ensure (letting the price being $+\infty$ 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 price-rates, i.e. the price 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 $x \in [0,1]$ (hence this guard is not displayed and transitions are only labeled by their respective discrete price): 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 valuation 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 (he 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 his 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

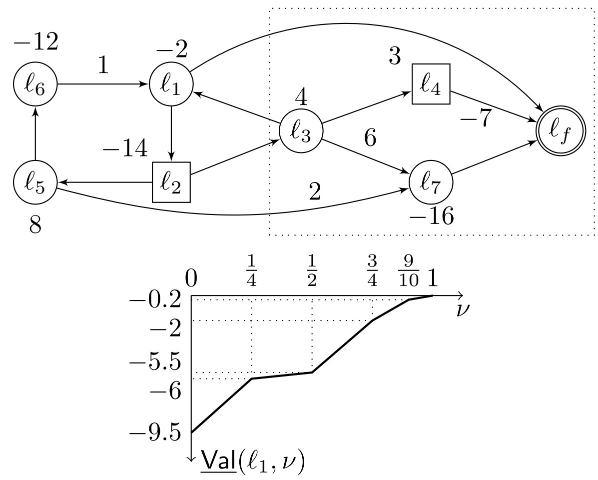

Figure 1: A simple priced timed game (left) and the lower value function of location ℓ_1 (right).

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 he will spend in ℓ_3 and whether he will go to ℓ_4 or ℓ_7. His optimal value from all (ℓ_3, ν) is thus $\inf_{0 \leqslant t \leqslant 1-\nu} \min\big(4t + (-3(\nu + t) - 4), 4t + 6 - 16(1 - (\nu + t))\big) = \min(-3\nu - 4, 16\nu - 10)$. Since $16\nu - 10 \geqslant -3\nu - 4$ if and only if $\nu \leqslant 6/19$, the best 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 subgame) is a piecewise affine function with two pieces.

Related work.

PTGs are a special case of hybrid games [dAHM01, MPS95, WT97], that were independently investigated in [BCFL04] and [ABM04]. For (non-necessarily turn-based) PTGs with *non-negative* prices, 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 valuation 0), is below a given threshold. They have also shown that, under the *strongly non-Zeno assumption* on prices (asking the existence of $\kappa > 0$ such that every cycle in the underlying region graph has a price 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 prices 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 prices and four or more clocks. In [BCJ09], the undecidability of the existence problem has also been shown for PTGs with arbitrary price-rates (without prices on transitions), and two or more clocks. On a positive side, the value problem was shown decidable by [BLMR06] for PTGs with one clock when the prices 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 prices on transitions, and price-rates restricted to two values amongst $\{-d, 0, +d\}$ (with $d \in \mathbf{N}$) was given in [BGK+14]. A survey summarising those results was done 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) prices. 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 modeling feature. Consider, for instance, a system which can consume but also produce energy at different rates. In this case, energy consumption could be modeled as a positive price-rate, and production by a negative price-rate. 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 an *exponential 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. 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. Thus, the whole class of PTGs with arbitrary weights and one clock 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 a price associated with the sequence of configurations and of transitions taken (hence, he is trying to minimise this price).

Note that this framework of quantitative reachability games that we develop here (and for which we prove a determinacy result, see Theorem 1) can be applied to other settings than our priced timed games. For example, special cases of quantitative reachability games are *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 paper, 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 fomalise 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 1 (Quantitative reachability games). *A* quantitative reachability game *is a tuple $G = (C = C_{\mathsf{Min}} \uplus C_{\mathsf{Max}}, \Sigma, E, F, p)$, where C is the set of* configurations *(that does not need to be finite, nor even countable), partitionned into the set C_{Min} of configurations of player* Min*, and the set C_{Max} of configurations of player* Max*; Σ is an alphabet whose elements are called letters; $E \subseteq C \times \Sigma \times C$ is the transition relation; $F \subseteq C$ is the set of* final configurations*; and $p\colon (C \times \Sigma)^* \times C \to \mathbf{R}$ maps each finite sequence $c_1 a_1 \cdots a_n c_n$ to a real number called the* price *of $c_1 a_1 \cdots a_n c_n$.*

For the sake of exposure, we assume that there are no deadlocks in the game, i.e. for all configurations $c \in C$, there exists $c' \in C$ and $a \in \Sigma$ such that $(c, a, c') \in E$. A *finite play* is a finite sequence $\rho = c_1 a_1 c_2 \dots c_n$ alternating between configurations and letters, and such that for all $i < n$: $(c_i, a_i, c_{i+1}) \in E$. In this case, we let $|\rho| = n$. A *play* is an infinite sequence $\rho = c_1 a_1 c_2 \cdots$ alternating between configurations and letters satisfying the same condition, i.e. for all i: $(c_i, a_i, c_{i+1}) \in E$. In that case, we let $|\rho|$ be the least position i such that $\ell_i \in L_f$, and $|\rho| = +\infty$ if there are no such positions. For the sake of clarity, we denote a play $c_1 a_1 c_2 \cdots$ as $c_1 \xrightarrow{a_1} c_2 \xrightarrow{a_2} \cdots$, and similarly for finite plays.

We take the viewpoint of player Min who wants to reach a final configuration. Thus, the *price* of a play $\rho = c_1 \xrightarrow{a_1} c_2 \cdots$, denoted $\mathsf{Price}(\rho)$ is either $+\infty$ if $|\rho| = +\infty$ (this is the worst situation for Min, which explains why the price is maximal in this case); or $p(c_1 \xrightarrow{a_1} c_2 \cdots c_n)$ if $|\rho| = n$.

A *strategy* for player Min is a function σ_{Min} mapping every finite play ending in a configuration $c \in C_{\mathsf{Min}}$ to a transition $(c, a, c') \in E$. Strategies σ_{Max} of player Max are defined accordingly. We let $\mathsf{Strat}_{\mathsf{Min}}(G)$ and $\mathsf{Strat}_{\mathsf{Max}}(G)$ be the sets of strategies of Min and Max, respectively. A pair $(\sigma_{\mathsf{Min}}, \sigma_{\mathsf{Max}}) \in \mathsf{Strat}_{\mathsf{Min}}(\mathcal{G}) \times \mathsf{Strat}_{\mathsf{Max}}(\mathcal{G})$ is called a *profile of strategies*. Together with an initial configuration c_0, it defines a unique play $\mathsf{Play}(c_0, \sigma_{\mathsf{Min}}, \sigma_{\mathsf{Max}}) = c_0 \xrightarrow{a_0} c_1 \cdots$ such that for all $i \geqslant 0$: $(c_i, a_i, c_{i+1}) = \sigma_{\mathsf{Min}}(c_0 \xrightarrow{a_0} \cdots c_i)$ if $c_i \in C_{\mathsf{Min}}$; and $(c_i, a_i, c_{i+1}) = \sigma_{\mathsf{Max}}(c_0 \xrightarrow{a_0} \cdots c_i)$ if $c_i \in C_{\mathsf{Max}}$. We let $\mathsf{Play}(\sigma_{\mathsf{Min}})$ (respectively, $\mathsf{Play}(c_0, \sigma_{\mathsf{Min}})$) be the set of plays that conform with σ_{Min} (and start in c_0), and define $\mathsf{Play}(\sigma_{\mathsf{Max}})$ and $\mathsf{Play}(c_0, \sigma_{\mathsf{Max}})$ accordingly. Given an initial configuration c_0, the price of a strategy σ_{Min} of Min is:

$$\mathsf{Price}(c_0, \sigma_{\mathsf{Min}}) = \sup_{\rho \in \mathsf{Play}(c_0, \sigma_{\mathsf{Min}})} \mathsf{Price}(\rho)\,.$$

It matches the intuition to be the largest price that Min may pay while following strategy σ_{Min}. This definition is equal to $\sup_{\sigma_{\mathsf{Max}}} \mathsf{Price}(\mathsf{Play}(c_0, \sigma_{\mathsf{Min}}, \sigma_{\mathsf{Max}}))$, 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, we define the price of σ_{Max} as:

$$\mathsf{Price}(c_0, \sigma_{\mathsf{Max}}) = \inf_{\rho \in \mathsf{Play}(c_0, \sigma_{\mathsf{Max}})} \mathsf{Price}(\rho) = \inf_{\sigma_{\mathsf{Min}}} \mathsf{Price}(\mathsf{Play}(c_0, \sigma_{\mathsf{Min}}, \sigma_{\mathsf{Max}}))\,.$$

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_0 arise, depending on which player chooses its strategy first. The *upper value* of c_0, defined as:

$$\overline{\mathsf{Val}}(c_0) = \inf_{\sigma_{\mathsf{Min}}} \sup_{\sigma_{\mathsf{Max}}} \mathsf{Price}(\mathsf{Play}(c_0, \sigma_{\mathsf{Min}}, \sigma_{\mathsf{Max}}))\,,$$

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

$$\underline{\mathsf{Val}}(c_0) = \sup_{\sigma_{\mathsf{Max}}} \inf_{\sigma_{\mathsf{Min}}} \mathsf{Price}(\mathsf{Play}(c_0, \sigma_{\mathsf{Min}}, \sigma_{\mathsf{Max}}))\,,$$

corresponds to the least price that Min can ensure when choosing its strategy *after* Max. It is easy to see that $\underline{\mathsf{Val}}(c_0) \leqslant \overline{\mathsf{Val}}(c_0)$, which explains the chosen names. Indeed, if Min picks its strategy after Max, he 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 his strategy first. This result, known as the *determinacy property*, is formalised here:

Theorem 1 (Determinacy of quantitative reachability games)**.** *For all quantitative reachability games G and configurations c_0, $\overline{\mathsf{Val}}(c_0) = \underline{\mathsf{Val}}(c_0)$.*

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 $G = (C = C_{\mathsf{Min}} \uplus C_{\mathsf{Max}}, \Sigma, E, F, p)$ to a family of such Gale-Stewart games $\mathit{Threshold}(G, r)$ parametrised by a threshold $r \in \mathbf{R}$.

The Gale-Stewart game $\mathit{Threshold}(G, r)$ 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.

In our case, the vertices of $\mathit{Threshold}(G, r)$ are the finite plays $c_0 \xrightarrow{a_0} c_1 \cdots c_n$ of G starting from configuration c_0. Such a vertex $v = c_0 \xrightarrow{a_0} c_1 \cdots c_n$ is owned by Min iff $c_n \in C_{\mathsf{Min}}$; otherwise v belongs to Max. A vertex $v = c_0 \xrightarrow{a_0} c_1 \cdots c_n$ has successors iff $c_n \notin F$. In this case, the successors of v are all the vertices $v \xrightarrow{a} c$ s.t. $(c_n, a, c) \in E$. Finally, a leaf $c_0 \xrightarrow{a_0} \cdots c_n$ (thus, with $c_n \in F$) is winning for Min iff $p(c_0 \xrightarrow{a_0} \cdots c_n) \leqslant r$.

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

$$\mathit{Win} = \bigcup_{v \in L \text{ s.t. } p(v) \leqslant r} \{\mathit{branch}(v)\}$$

where L is the set of leaves of $\mathit{Threshold}(G, r)$, and $\mathit{branch}(v)$ is the (unique) branch from c_0 to v. Then,

$$\mathit{Win} = \bigcup_{v \in L \text{ s.t. } p(v) \leqslant r} \mathit{Cone}(v)$$

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 $r \in \mathbf{R}$ i.e. either Min or Max has a winning strategy from the root of the tree. Notice that strategies in G and $Threshold(G, r)$ are isomorphic.

We rely on this result to prove that $\underline{\mathsf{Val}}(c_0) \geqslant \overline{\mathsf{Val}}(c_0)$ in G (the other inequality being always true). We consider two cases:

1. If $\overline{\mathsf{Val}}(c_0) = -\infty$, then $\underline{\mathsf{Val}}(c_0)$ being at most $\overline{\mathsf{Val}}(c_0)$ is $-\infty$ too.
2. If $\overline{\mathsf{Val}}(c_0) > -\infty$, consider any real number t such that $t < \overline{\mathsf{Val}}(c_0)$. By definition of the upper value, for all strategies σ_{Min}, we have $\mathsf{Price}(c_0, \sigma_{\mathsf{Min}}) > t$. Therefore, Min loses in the game $Threshold(G, t)$. By determinacy, Max wins in this game, i.e. there exists a strategy σ^t_{Max} such that $\mathsf{Price}(c_0, \sigma^t_{\mathsf{Max}}) > t$. By definition of the lower value, this ensures that $\underline{\mathsf{Val}}(c_0) \geqslant t$. Therefore, $t < \overline{\mathsf{Val}}(c_0)$ implies $t \leqslant \underline{\mathsf{Val}}(c_0)$: since this holds for all t, we have $\underline{\mathsf{Val}}(c_0) \leqslant \overline{\mathsf{Val}}(c_0)$. □

In such determined games, we denote by Val the value of the game, defined as $\mathsf{Val} = \overline{\mathsf{Val}} = \underline{\mathsf{Val}}$.

3. Priced timed games

We are now ready to formally introduce the core model of our paper: 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*, see Section 3.4) that we will rely upon in our algorithms to solve priced timed games.

3.1. Notations and definitions.

Let x denote a positive real-valued variable called *clock*. A *guard* (or *clock constraint*) is an interval with endpoints in $\mathbf{N} \cup \{+\infty\}$. We often abbreviate guards, writing for instance $x \leqslant 5$ instead of $[0, 5]$. Let $S \subseteq \mathsf{Guard}(x)$ be a finite set of guards. We let $[\![S]\!] = \bigcup_{I \in S} I$. Assuming $M_0 = 0 < M_1 < \cdots < M_k$ are all the endpoints of the intervals in S (to which we add 0 if needed), we let $\mathsf{Reg}_S = \{(M_i, M_{i+1}) \mid 0 \leqslant i \leqslant k-1\} \cup \{\{M_i\} \mid 0 \leqslant i \leqslant k\}$ be the set of *regions* of 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 $S \subseteq \mathsf{Guard}(x)$, a *cost function* over S is a function $f : [\![\mathsf{Reg}_S]\!] \to \overline{\mathbf{R}} = \mathbf{R} \cup \{+\infty, -\infty\}$ such that over all regions $r \in \mathsf{Reg}_S$, f is either infinite or a continuous piecewise affine function, with a finite set of cutpoints (points where the first derivative is not defined) $\{\kappa_1, \dots, \kappa_p\} \subseteq \mathbf{Q}$, and with $f(\kappa_i) \in \mathbf{Q}$ for all $1 \leqslant i \leqslant p$. 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 $f \in \mathsf{CF}_S$ and $f' \in \mathsf{CF}_{S'}$ be two costs functions on set of guards $S, S' \subseteq \mathsf{Guard}(x)$, such that $[\![S]\!] \cap [\![S']\!]$ is a singleton. We let $f \rhd f'$ be the cost function in $\mathsf{CF}_{S\cup S'}$ such that $(f \rhd f')(\nu) = f(\nu)$ for all $\nu \in [\![\mathsf{Reg}_S]\!]$, and $(f \rhd f')(\nu) = f'(\nu)$ for all $\nu \in [\![\mathsf{Reg}_{S'}]\!] \setminus [\![\mathsf{Reg}_S]\!]$. For example, let $S = [0, 1)$ and $S' = S \cup \{1\}$. We define the cost functions f_1 and f_2 such that f_1 is equal to $+\infty$ on the set of regions $\mathsf{Reg}_{S'}$ and f_2 is equal to 0 on the set of regions Reg_S. The cost function $f_2 \rhd f_1 \in \mathsf{CF}_{S'}$ is equal to 0 on $[0, 1)$ and to $+\infty$ on $\{1\}$ and the cost function $f_1 \rhd f_2 \in \mathsf{CF}_{S'}$ is equal to $+\infty$ on $[0, 1]$. Thus $f_1 \rhd f_2$ is equal to f_1 while $f_2 \rhd f_1$ extends f_2 with a $+\infty$ value in $\{1\}$.

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 price to be paid when ending the game in this location.

Definition 2. *A* priced timed game *(PTG for short) $\mathcal{G}$ is a tuple* $(L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_f, L_u, \boldsymbol{\varphi}, \Delta, \pi)$ *where:*

- L_{Min} *and* L_{Max} *are finite sets of* locations *belonging respectively to player* Min *and* Max*. We assume* $L_{\mathsf{Min}} \cap L_{\mathsf{Max}} = \emptyset$*;*
- L_f *is a finite set of* final *locations. We let* $L = L_{\mathsf{Min}} \cup L_{\mathsf{Max}} \cup L_f$ *be the set of all locations of the PTG;*
- $L_u \subseteq L \setminus L_f$ *is the set of* urgent *locations*[2]*;*
- $\Delta \subseteq (L \setminus L_f) \times \mathsf{Guard}(x) \times \{\top, \bot\} \times L$ *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;*
- $\boldsymbol{\varphi} = (\varphi_\ell)_{\ell \in L_f}$ *associates to all locations* $\ell \in L_f$ *a* final cost function*, that is an affine*[3] *cost function* φ_ℓ*;*
- $\pi \colon L \cup \Delta \to \mathbf{Z}$ *is a mapping associating an integer price to all locations and transitions. In the case of a location* ℓ*, we say that* $\pi(\ell)$ *is* ℓ*'s* price-rate*.*

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 $M \in \mathbf{N}$ such that for all guards $I \in S_{\mathcal{G}}$, $I \subseteq [0, M]$.[4] We denote by $\mathsf{Reg}_{\mathcal{G}}$ the set $\mathsf{Reg}_{S_{\mathcal{G}}}$ of *regions of* $\mathcal{G}$. We further denote[5] by $\Pi^{\mathrm{tr}}_{\mathcal{G}}$, $\Pi^{\mathrm{loc}}_{\mathcal{G}}$ and $\Pi^{\mathrm{fin}}_{\mathcal{G}}$ respectively the values $\max_{\delta\in\Delta} |\pi(\delta)|$, $\max_{\ell\in L} |\pi(\ell)|$ and $\sup_{\nu\in[0,M]} \max_{\ell\in L} |\varphi_\ell(\nu)| = \max_{\ell\in L} \max(|\varphi_\ell(0)|, |\varphi_\ell(M)|)$ (the last equality holds because we have assumed that φ_ℓ is affine). That is, $\Pi^{\mathrm{tr}}_{\mathcal{G}}$, $\Pi^{\mathrm{loc}}_{\mathcal{G}}$ and $\Pi^{\mathrm{fin}}_{\mathcal{G}}$ are the largest absolute values of the location prices, transition prices and final cost functions.

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

[3]Not that in our setting, an affine function is of the form $f(\nu) = a \times \nu + b$.

[4]Observe that this 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 prices.

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

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

$$G_{\mathcal{G}} = \big(\mathsf{Conf}_{\mathcal{G}}, \Sigma = (\mathbf{R}^+ \times \Delta \times \mathbf{R}), E, F = (L_f \times \mathbf{R}^+), p\big)$$

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 Price, $\overline{\mathsf{Val}}$, $\underline{\mathsf{Val}}$ and Val functions, and the notions of plays from $G_{\mathcal{G}}$ to $\mathcal{G}$. A *configuration* of $\mathcal{G}$ is a pair $s = (\ell, \nu) \in L \times \mathbf{R}^+$, where ℓ and ν are respectively the current location and clock values of $\mathcal{G}$. We denote by $\mathsf{Conf}_{\mathcal{G}}$ 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 $t \in \mathbf{R}^+$ be a delay. Then, $((\ell, \nu), (t, \delta, c), (\ell', \nu')) \in E$, iff:

(i) $\ell \in L_u$ implies $t = 0$ (no time can elapse in urgent locations);
(ii) $\nu + t \in I$ (the guard is satisfied);
(iii) $R = \top$ implies $\nu' = 0$ (when the clock is reset);
(iv) $R = \bot$ implies $\nu' = \nu + t$ (when the clock is not reset);
(v) $c = \pi(\delta) + t \times \pi(\ell)$ (the price of (t, δ) takes into account the price-rate of ℓ, the delay spent in ℓ, and the price of δ).

In this case, we say that there is a (t, δ)-transition from (ℓ, ν) to (ℓ', ν') with price 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 $\mathbf{R}^+ \times \Delta \times \mathbf{R}$, and its set of configurations is $\mathsf{Conf}_{\mathcal{G}}$, 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 price function p is obtained by summing the price of the play (transitions and time spent in the locations) and the final cost function if applicable. Formally, for a finite play $\rho = (\ell_1, \nu_1) \xrightarrow{t_1,\delta_1,c_1} (\ell_2, \nu_2) \cdots (\ell_n, \nu_n)$ with $\forall k < n, \ell_k \notin L_f$, if $\ell_n \in L_f$ then $p(\rho) = \sum_{i=1}^{n-1} c_i + \varphi_{\ell_{|\rho|}}(\nu_n)$ else $p(\rho) = \sum_{i=1}^{n-1} c_i$.

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 1, 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 $\mathsf{Val}_{\mathcal{G}}$ 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 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 $t_{n,m} = (\ell_n, [0,1], \bot, \ell_m)$ ends in the unique final location ℓ_f and its price is $p(\rho) = 0 - 3.5 + 6 - 12 = -9.5$.

Let us fix a PTG $\mathcal{G}$ with initial configuration c_0. We say that a strategy σ_{Min} of Min is *optimal* if $\mathsf{Price}(c_0, \sigma_{\mathsf{Min}}) = \mathsf{Val}_{\mathcal{G}}(c_0)$, i.e., it ensures Min to enforce the value of the game, whatever Max does. Similarly, σ_{Min} is *ε-optimal* if $\mathsf{Price}(c_0, \sigma_{\mathsf{Min}}) \leqslant \mathsf{Val}_{\mathcal{G}}(c_0) + \varepsilon$. And, symmetrically, a strategy σ_{Max} of Max is *optimal* (respectively, *ε-optimal*) if $\mathsf{Price}(c_0, \sigma_{\mathsf{Max}}) = \mathsf{Val}_{\mathcal{G}}(c_0)$ (respectively, $\mathsf{Price}(c_0, \sigma_{\mathsf{Max}}) \geqslant \mathsf{Val}_{\mathcal{G}}(c_0) + \varepsilon$).

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 your knowledge, original) result. For all locations ℓ, let $\mathsf{Val}_{\mathcal{G}}(\ell)$ denote the function such that $\mathsf{Val}_{\mathcal{G}}(\ell)(\nu) = \mathsf{Val}_{\mathcal{G}}(\ell,\nu)$ for all $\nu \in \mathbf{R}^+$. Then, we show that, for all ℓ, $\mathsf{Val}_{\mathcal{G}}(\ell)$ is a *piecewise continuous function* that might exhibit discontinuities *only on the borders of the regions* of $\mathsf{Reg}_{\mathcal{G}}$.

Theorem 2. *For all (one-clock) PTGs $\mathcal{G}$, for all $r \in \mathsf{Reg}_{\mathcal{G}}$, for all $\ell \in L$, $\mathsf{Val}_{\mathcal{G}}(\ell)$ is either infinite or continuous over r.*

Proof. Our goal is to show that for every location ℓ, region $r \in \mathsf{Reg}_{\mathcal{G}}$ and valuations ν and ν' in r,

$$|\mathsf{Val}(\ell,\nu) - \mathsf{Val}(\ell,\nu')| \leqslant \Pi^{\mathrm{loc}}|\nu - \nu'|.$$

This is equivalent to showing:

$$\mathsf{Val}(\ell,\nu) \leqslant \mathsf{Val}(\ell,\nu') + \Pi^{\mathrm{loc}}|\nu-\nu'| \quad \text{and} \quad \mathsf{Val}(\ell,\nu') \leqslant \mathsf{Val}(\ell,\nu) + \Pi^{\mathrm{loc}}|\nu-\nu'|\,.$$

As those two equations are symmetric with respect to ν and ν', we only have to show either of them. We will 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 $\mathsf{Price}((\ell,\nu'),\sigma'_{\mathsf{Min}}) \leqslant \mathsf{Price}((\ell,\nu),\sigma_{\mathsf{Min}}) + \Pi^{\mathrm{loc}}|\nu-\nu'|$. 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 $\rho' = \mathsf{Play}((\ell,\nu'),\sigma'_{\mathsf{Min}},\sigma_{\mathsf{Max}})$ for some strategy σ_{Max} of Max) to plays from (ℓ,ν), consistent with σ_{Min}, such that:

$$\mathsf{Price}(\rho') \leqslant \mathsf{Price}(g(\rho')) + \Pi^{\mathrm{loc}}|\nu-\nu'|\,.$$

Let $r \in \mathsf{Reg}_{\mathcal{G}}$, $\nu,\nu' \in r$ and σ_{Min} be a strategy of Min. We define σ'_{Min} and g by induction on the size of their arguments; more precisely, we define $\sigma'_{\mathsf{Min}}(\rho'_1)$ and $g(\rho'_2)$ by induction on k, for all plays ρ'_1 and ρ'_2 from (ℓ,ν'), consistent with σ'_{Min} of size $k-1$ and k, respectively. We also show during this induction that for each play $\rho' = (\ell_0,\nu'_0) \xrightarrow{c'_0} \cdots \xrightarrow{c'_{k-1}} (\ell_k,\nu'_k)$ from (ℓ,ν'), consistent with σ'_{Min}, if we let $(\ell_0,\nu_0) \xrightarrow{c_0} \cdots \xrightarrow{c_{\ell-1}} (\ell_\ell,\nu_\ell) = g(\rho')$:

(i) ρ' and $g(\rho')$ have the same length, i.e. $|\rho| = \ell = k = |\rho'|$,
(ii) for every $i \in \{1,\dots,k\}$, ν_i and ν'_i are in the same region, i.e. there exists a region $r' \in \mathsf{Reg}_{\mathcal{G}}$ such that $\nu_i \in r'$ and $\nu'_i \in r'$,
(iii) $|\nu_k - \nu'_k| \leqslant |\nu-\nu'|$,
(iv) $\mathsf{Price}(\rho') \leqslant \mathsf{Price}(g(\rho')) + \Pi^{\mathrm{loc}}(|\nu-\nu'| - |\nu_k-\nu'_k|)$.

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

If $k=0$, σ'_{Min} does not have to be defined. Moreover, in that case, $\rho' = (\ell,\nu')$ and $g(\rho') = (\ell,\nu)$. Both plays have size 0, ν and ν' are in the same region by hypothesis of the lemma, and $\mathsf{Price}(\rho') = \mathsf{Price}(g(\rho')) = 0$, therefore all four properties are true.

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

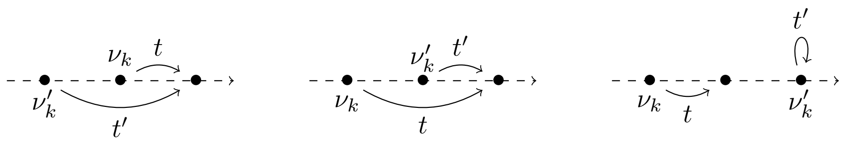

(a) (b) (c)

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

consider a play $\rho' = (\ell_1, \nu'_1) \xrightarrow{c'_1} \cdots \xrightarrow{c'_{k-1}} (\ell_k, \nu'_k)$ from (ℓ, ν'), consistent with σ'_{Min} 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. $\sigma_{\mathsf{Min}}(g(\rho')) = (t, \delta)$. Then, we define $\sigma'_{\mathsf{Min}}(\rho') = (t', \delta)$ where $t' = \max(0, \nu_k + t - \nu'_k)$. The delay t' respects the guard of transition δ since either $\nu_k + t = \nu'_k + t'$ or $\nu_k \leqslant \nu_k + t \leqslant \nu'_k$, 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. This is illustrated in Figure 2.

Let us now build the mapping g. Let $\rho' = (\ell_1, \nu'_1) \xrightarrow{c'_1} \cdots \xrightarrow{c'_k} (\ell_{k+1}, \nu'_{k+1})$ be a play from (ℓ, ν') consistent with σ'_{Min} and $\tilde{\rho}' = (\ell_1, \nu'_1) \xrightarrow{c'_1} \cdots \xrightarrow{c'_{k-1}} (\ell_k, \nu'_k)$ its prefix of size 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 $g(\tilde{\rho}') = (\ell_1, \nu_1) \xrightarrow{c_1} \cdots \xrightarrow{c_{k-1}} (\ell_k, \nu_k)$ (with $(\ell_1, \nu_1) = (\ell, \nu)$) verifies properties i, ii, iii and iv. If ℓ_k is a location of Min and $\sigma_{\mathsf{Min}}(g(\tilde{\rho}')) = (t, \delta)$, 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 valuation 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: i 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 ii. To prove iii, notice that we always have either $\nu_k + t = \nu'_k + t'$ or $\nu_k \leqslant \nu_k + t \leqslant \nu'_k = \nu'_k + t'$ or $\nu'_k \leqslant \nu'_k + t \leqslant \nu_k = \nu_k + t$. In all of these possibilities, we have $|(\nu_k + t) - (\nu'_k + t')| \leqslant |\nu_k - \nu'_k|$. We finally check property iv. In both cases:

$$\begin{aligned}\mathsf{Price}(\rho') &= \mathsf{Price}(\tilde{\rho}') + \pi(\delta) + t'\pi(\ell_k)\\ &\leqslant \mathsf{Price}(g(\tilde{\rho}')) + \Pi^{\mathrm{loc}}(|\nu - \nu'| - |\nu_k - \nu'_k|) + \pi(\delta) + t'\pi(\ell_k)\\ &= \mathsf{Price}(g(\rho')) + (t' - t)\pi(\ell_k) + \Pi^{\mathrm{loc}}(|\nu - \nu'| - |\nu_k - \nu'_k|)\,.\end{aligned}$$

If δ contains no reset, let us prove that

$$|t' - t| = |\nu_k - \nu'_k| - |\nu'_{k+1} - \nu_{k+1}|\,. \tag{1}$$

Indeed, since $t' = \nu'_{k+1} - \nu'_k$ 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 $t' + \nu'_k = t + \nu_k$, $\nu'_{k+1} = \nu_{k+1}$, thus $|t' - t| = |\nu_k - \nu'_k| = |\nu_k - \nu'_k| - |\nu'_{k+1} - \nu_{k+1}|$.

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

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, either $t = 0$ and $t' \leqslant \nu_k - \nu'_k$, or $t' = 0$ and $t \leqslant \nu'_k - \nu_k$.

In all cases, we have proved (1). Coupled with the fact that $|P(\ell_k)| \leqslant \Pi^{\mathrm{loc}}$, we conclude that:

$$\mathsf{Price}(\rho') \leqslant \mathsf{Price}(g(\rho')) + \Pi^{\mathrm{loc}}(|\nu - \nu'| - |\nu_{k+1} - \nu'_{k+1}|)\,.$$

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 valuation of its last configuration, and ν_k be the clock valuation of the last configuration of $g(\rho')$. Combining (*iii*) and (*iv*) we have $\mathsf{Price}(\rho') \leqslant \mathsf{Price}(g(\rho')) + \Pi^{\mathrm{loc}}|\nu - \nu'|$ which concludes the proof. $\square$

Remark 1. *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 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 price of this cycle is null, and that no other delays can be played, hence $\overline{\mathsf{Val}}(\ell_0, 0, 0.5) = 0$. However, starting from a configuration $(\ell_0, 0, 0.6)$, and following the same path, yields the cycle $(\ell_0, 0, 0.6) \xrightarrow{0,e_0,0} (\ell_1, 0, 0.6) \xrightarrow{0.4,e_1,2} (\ell_2, 0.4, 0) \xrightarrow{0.6,e_2,-3} (\ell_0, 0, 0.6)$ with price -1. Hence, $\overline{\mathsf{Val}}(\ell_0, 0, 0.6) = -\infty$, and the function is not continuous although both valuations $(0, 0.5)$ 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.*

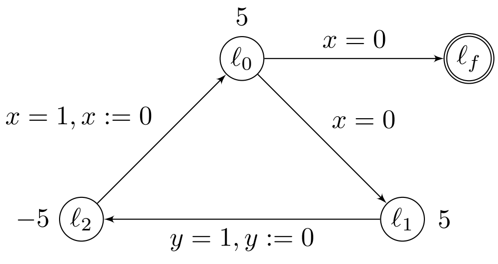

Figure 3: A PTG with 2 clocks whose value function is not continuous inside a region.

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 prices. Formally, an r-SPTG, with $r \in \mathbf{Q}^+ \cap [0,1]$, is a PTG $\mathcal{G} = (L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_f, L_u, \boldsymbol{\varphi}, \Delta, \pi)$ such that for all transitions $(\ell, I, R, \ell') \in \Delta$, $I = [0, r]$ (the clock is also bounded by r) and $R = \bot$. 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. *Let $\mathcal{G}$ be an SPTG. Then, for all locations $\ell \in L$, the function $\mathsf{Val}(\ell)$ is either infinite, or continuous and piecewise-affine with at most an exponential number of cutpoints. The value functions for all locations, as well as a pair of optimal strategies $(\sigma_{\mathsf{Min}}, \sigma_{\mathsf{Max}})$ (that always exist if no values are infinite) can be computed in exponential time.*

3.4. Switching strategies.

Let us now discuss a class of (simple) strategies that are sufficient to play optimally. Those strategies, called *switching strategies*, will be instrumental in proving the result above. Roughly speaking, Max has always a *memoryless* optimal strategy, while Min might need *(finite) memory* to play optimally—it is already the case in untimed quantitative reachability games with arbitrary weights [BGHM16]. Moreover, these strategies are finitely representable (recall that even a memoryless strategy depends on the current *configuration* and that there are infinitely many in our time setting).

We start by formalising Max's strategies, thanks to the notion of *finite positional strategies*:

Definition 3 (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 $0 \leqslant \nu_1^\ell < \nu_2^\ell < \cdots < \nu_k^\ell = 1$ and a finite sequence of transitions $\delta_1, \ldots, \delta_k \in \Delta$ such that*

(i) *for all $1 \leqslant i \leqslant k$, for all $\nu \in (\nu_{i-1}^\ell, \nu_i^\ell]$, either $\sigma(\ell, \nu) = (0, \delta_i)$, or $\sigma(q, \nu) = (\nu_i^\ell - \nu, \delta_i)$ (assuming $\nu_0^\ell = \min(0, \nu_1^\ell)$); and*

(ii) *if $\nu_1^\ell > 0$, then $\sigma(\ell, 0) = (\nu_1^\ell, \delta_1)$.*

We let $\mathsf{pts}(\sigma)$ be the set of ν_i^ℓ for all ℓ and i, and $\mathsf{int}(\sigma)$ be the set of all successive intervals generated by $\mathsf{pts}(\sigma)$. Finally, we let $|\sigma| = |\mathsf{int}(\sigma)|$ be the size of σ. Intuitively, in an interval $(\nu_{i-1}^\ell, \nu_i^\ell]$, σ always returns the same move: either to take *immediately* δ_i or to wait until the clock reaches the endpoint ν_i^ℓ and then take δ_i.

While Max's strategies can be memoryless, we observe that Min may require memory to play optimally as shown in the following example taken from [BGHM16]. Consider the SPTG of Figure 4, where W is a positive integer, and every location has price-rate 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 either

(i) transition (ℓ_1, ℓ_3) (with weight $-W$) will eventually be traversed; or

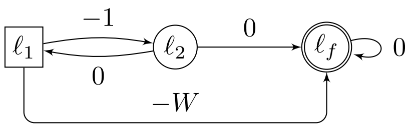

Figure 4: An SPTG where Min needs memory to play optimally

(ii) 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 he cannot ensure a lower price, since Max always has the opportunity to take the transition (ℓ_1, ℓ_f) (with weight $-W$) 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 Min strategy we have just given is optimal, and there is *no optimal memoryless strategy* 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 $+\infty$.

Informally, we will compute optimal *switching* strategies, as introduced in [BGHM16] (in the untimed setting).

Definition 4 (Switching strategies)**.** *A switching strategy is described by a pair $(\sigma^1_{\mathsf{Min}}, \sigma^2_{\mathsf{Min}})$ of FP-strategies and a switch threshold K, and consists in playing σ^1_{Min} until the total accumulated price of the discrete transitions is below K; and then to* switch *to strategy σ^2_{Min}.*

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

Formally, an NC-strategy σ_{Min} of Min is an FP-strategy such that for all runs $\rho = (\ell_1, \nu) \xrightarrow{c_1} \cdots \xrightarrow{c_{k-1}} (\ell_k, \nu') \in \mathsf{Play}(\sigma_{\mathsf{Min}})$ with $\ell_1 = \ell_k$, and ν, ν' in the same interval of $\mathsf{int}(\sigma_{\mathsf{Min}})$, the sum of prices of *discrete transitions* is at most -1, i.e. $\pi(\ell_1, \ell_2) + \cdots + \pi(\ell_{k-1}, \ell_k) \leqslant -1$. To characterise the fact that σ_{Min} must allow Min to reach a price which is *small enough, without necessarily reaching a target state*, we define the *fake value* of an NC-strategy σ_{Min} from a configuration s as $\mathsf{fake}^{\sigma_{\mathsf{Min}}}_{\mathcal{G}}(s) = \sup\{\mathsf{Price}(\rho) \mid \rho \in \mathsf{Play}(s, \sigma_{\mathsf{Min}}), \rho \text{ reaches a target}\}$, i.e. the value obtained when *ignoring* the σ_{Min}-induced plays that *do not* reach the target. Thus, clearly, $\mathsf{fake}^{\sigma_{\mathsf{Min}}}_{\mathcal{G}}(s) \leqslant \mathsf{Val}^{\sigma_{\mathsf{Min}}}(s)$. We say that an NC-strategy is *fake-optimal* if its fake value, in every configuration, is equal to the optimal value of the configuration in the game. This is justified by the following result whose proof relies on the switching strategies described before:

Lemma 4. *If $\mathsf{Val}_{\mathcal{G}}(\ell, \nu) \neq +\infty$, for all ℓ and ν, then for all NC-strategies σ_{Min}, there is a strategy σ'_{Min} such that $\mathsf{Val}^{\sigma'_{\mathsf{Min}}}_{\mathcal{G}}(s) \leqslant \mathsf{fake}^{\sigma_{\mathsf{Min}}}_{\mathcal{G}}(s)$ for all configurations*

s. *In particular, if* σ_{Min} *is a fake-optimal NC-strategy, then* σ'_{Min} *is an optimal (switching) strategy of the SPTG.*

Proof. First of all, notice that all finite plays $\rho \in \mathsf{Play}(\sigma_{\mathsf{Min}})$ with all clock valuations in the same interval I of $\mathsf{int}(\sigma)$ verify $\mathsf{Price}(\rho) \leqslant |I|\Pi^{\mathrm{loc}} + |L|\Pi^{\mathrm{tr}} - |\rho|/|L|$. Indeed, the price of ρ is the sum of the price generated by staying in locations, which is bounded by $|I|\Pi^{\mathrm{loc}}$, and the price of the transitions. One can extract at least $|\rho|/|L|$ cycles with transition prices at most -1 (by definition of an NC-strategy), and what remains is of size at most $|L|$, ensuring that the transition price is bounded by $|L|\Pi^{\mathrm{tr}} - |\rho|/|L|$.

Then, by splitting runs among intervals of $\mathsf{int}(\sigma_{\mathsf{Min}})$, we can easily obtain that all finite plays $\rho \in \mathsf{Play}(\sigma_{\mathsf{Min}})$ verify $\mathsf{Price}(\rho) \leqslant \Pi^{\mathrm{loc}} + (2|\sigma_{\mathsf{Min}}|-1) \times |L|\Pi^{\mathrm{tr}} - (|\rho| - |\sigma_{\mathsf{Min}}|)/|L|$. Indeed, letting $I_1, I_2, \ldots, I_k$ the interval of $\mathsf{int}(\sigma_{\mathsf{Min}})$ visited during ρ (with $k \leqslant |\sigma_{\mathsf{Min}}|$), one can split ρ into k runs $\rho = \rho_1 \xrightarrow{c_1} \rho_2 \xrightarrow{c_2} \cdots \rho_k$ such that in ρ_i all clock values are in I_i (remember that SPTGs contain no reset transitions). By the previous inequality, we have $\mathsf{Price}(\rho_i) \leqslant |I_i|\Pi^{\mathrm{loc}} + |L|\Pi^{\mathrm{tr}} - |\rho_i|/|L|$. Thus, also splitting prices c_i with respect to discrete price and price of delaying, we obtain $\mathsf{Price}(\rho) = \sum_{i=1}^{k} \mathsf{Price}(\rho_i) + \sum_{i=1}^{k-1} c_i \leqslant (2|\sigma_{\mathsf{Min}}| - 1) \times |L|\Pi^{\mathrm{tr}} + \Pi^{\mathrm{loc}} - (|\rho| - |\sigma_{\mathsf{Min}}|)/|L|$, since $|\rho| \leqslant \sum_i |\rho_i| + k \leqslant \sum_i |\rho_i| + |\sigma_{\mathsf{Min}}|$ and $\sum_i |I_i| \leqslant 1$.

We now turn to the proof of the lemma. To that extent, we suppose known an attractor strategy for Min, i.e. a strategy that ensures to reach a final location: it exists thanks to the hypothesis on the finiteness of the values. From every configuration, it reaches a final location with a price bounded above by a given constant M. Notice first that, under the hypothesis that no configurations of the SPTG have value $-\infty$, we have $\mathsf{fake}_{\mathcal{G}}^{\sigma_{\mathsf{Min}}}(s) > -\infty$ for a configuration s. Otherwise, consider the strategy σ'_{Min} obtained by playing σ_{Min} until having computed a price bounded above by a fixed integer $N \in \mathbf{Z}$, in which case we switch to the attractor strategy. By the previous inequality, the switch is sure to happen since the right term tends to $-\infty$ when the length of ρ tends to ∞. Then, we know that the value guaranteed by σ'_{Min} is at most N, implying that the optimal value $\mathsf{Val}(s)$ is $-\infty$, which contradicts the hypothesis. Then, to prove the result of the lemma, consider the strategy σ'_{Min} obtained by playing σ_{Min} until having computed a price bounded above by the finite value $\mathsf{fake}_{\mathcal{G}}^{\sigma_{\mathsf{Min}}}(s) - M$, in which case we switch to the attractor strategy. Once again, the switch is sure to happen, implying that every play conforming to σ_{Min} reaches the target: moreover, the price of such a play is necessarily at most $\mathsf{fake}^{\sigma_{\mathsf{Min}}}(s)$ by construction. Then, we directly obtain that $\mathsf{Val}_{\mathcal{G}}^{\sigma'_{\mathsf{Min}}}(s) \leqslant \mathsf{fake}_{\mathcal{G}}^{\sigma_{\mathsf{Min}}}(s)$. $\square$

Then, an SPTG is called *finitely optimal* if

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

The central point in establishing Theorem 3 will thus be to prove that **all SPTGs are finitely optimal**, as this guarantees the existence of well-behaved optimal strategies and value functions. We will also show that they can be computed in exponential time. The proof is by induction on the number of 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 $\mathsf{Val}_{\mathcal{G}}(\ell, 1)$ for all ℓ and all games $\mathcal{G}$ (even with non-urgent locations) since time cannot elapse anymore when the clock has valuation 1. Next, using the continuity result of Theorem 2, we can detect locations ℓ where $\mathsf{Val}_{\mathcal{G}}(\ell, \nu) \in \{+\infty, -\infty\}$, 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 prices).

4. SPTGs with only urgent locations

Throughout this section, we consider an r-SPTG $\mathcal{G} = (L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_f, L_u, \boldsymbol{\varphi}, \Delta, \pi)$ where all locations are urgent, i.e. $L_u = L_{\mathsf{Min}} \cup L_{\mathsf{Max}}$. Since all locations in $\mathcal{G}$ are urgent, we may extract from a play $\rho = (\ell_0, \nu) \xrightarrow{c_0} (\ell_1, \nu) \xrightarrow{c_1} \cdots$ the clock valuations, as well as prices $c_i = \pi(\ell_i, \ell_{i+1})$, hence denoting plays by their sequence of locations $\ell_0 \ell_1 \cdots$. The price of this play is $\mathsf{Price}(\rho) = +\infty$ if $\ell_k \notin L_f$ for all $k \geqslant 0$; and $\mathsf{Price}(\rho) = \sum_{i=0}^{k-1} \pi(\ell_i, \ell_{i+1}) + \varphi_{\ell_k}(\nu)$ if k is the least position such that $\ell_k \in L_f$.

4.1. Computing the value for a particular valuation

We first explain how we can compute the value function of the game for a *fixed* clock valuation $\nu \in [0, r]$: more precisely, we will compute the vector $(\mathsf{Val}(\ell, \nu))_{\ell \in L}$ of values for all locations. We will denote by $\mathsf{Val}_\nu(\ell)$ the value $\mathsf{Val}(\ell, \nu)$, 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 prices 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 $\mathsf{Val}_\nu(\ell) = +\infty$ and $\mathsf{Val}_\nu(\ell) = -\infty$ 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 price is paid only on the target location, all other prices being integers), for all plays ρ, $\mathsf{Price}(\rho)$ is a value from the set $\mathbf{Z}_{\nu,\boldsymbol{\varphi}} = \mathbf{Z} + \{\varphi_\ell(\nu) \mid \ell \in L_f\}$. We further define $\mathbf{Z}_{\nu,\boldsymbol{\varphi}}^{+\infty} = \mathbf{Z}_{\nu,\boldsymbol{\varphi}} \cup \{+\infty\}$. Clearly, $\mathbf{Z}_{\nu,\boldsymbol{\varphi}}$ contains at most $|L_f|$ values between two consecutive integers, i.e.

$$\forall i \in \mathbf{Z} \quad |[i, i+1] \cap \mathbf{Z}_{\nu,\boldsymbol{\varphi}}| \leqslant |L_f| \tag{2}$$

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

$$\mathcal{F}(\boldsymbol{x})_\ell = \begin{cases} \varphi_\ell(\nu) & \text{if } \ell \in L_f \\ \max\limits_{(\ell,\ell') \in \Delta} \big(\pi(\ell, \ell') + x_{\ell'}\big) & \text{if } \ell \in L_{\mathsf{Max}} \\ \min\limits_{(\ell,\ell') \in \Delta} \big(\pi(\ell, \ell') + x_{\ell'}\big) & \text{if } \ell \in L_{\mathsf{Min}}\,. \end{cases}$$

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

The intuition behind is that $\boldsymbol{x}^{(i)}$ *is the value of the game (when the clock takes valuation 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 $\rho = \ell_0\ell_1\cdots$, we let $\mathsf{Price}^{\leqslant i}(\rho) = \mathsf{Price}(\rho)$ if $\ell_k \in L_f$ for some $k \leqslant i$, and $\mathsf{Price}^{\leqslant i}(\rho) = +\infty$ otherwise. We further let

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

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

Lemma 5. *For all $i \geqslant 0$, and $\ell \in L$: $\boldsymbol{x}_\ell^{(i)} = \overline{\mathsf{Val}}_\nu^{\leqslant i}(\ell)$.*

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 $\mathcal{F}$. ☐

Now, let us study how the sequence $(\overline{\mathsf{Val}}_\nu^{\leqslant i})_{i\geqslant 0}$ 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 $(\mathbf{Z}_{\nu,\boldsymbol{\varphi}}^{+\infty})^L$), the sequence $(\overline{\mathsf{Val}}_\nu^{\leqslant i})_{i\geqslant 0}$ converges towards the greatest fixed point of $\mathcal{F}$. Let us now show that Val_ν is actually this greatest fixed point. First, Corollary 8 of [BGHM16] can be adapted to obtain

Lemma 6. *For all $\ell \in L$: $\overline{\mathsf{Val}}_\nu^{\leqslant |L|}(\ell) \leqslant |L|\Pi^{\mathrm{tr}} + \Pi^{\mathrm{fin}}$.*

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

(i) $j < k$ implies $\overline{\mathsf{Val}}_\nu^{\leqslant j}(\ell) = +\infty$ and

(ii) $j \geqslant k$ implies $\overline{\mathsf{Val}}_\nu^{\leqslant j}(\ell) \leqslant jW + \Pi^{\mathrm{fin}}$ and $\overline{\mathsf{Val}}_\nu^{\leqslant j}(\ell) \in \mathbf{Z}_{\nu,\boldsymbol{\varphi}}$.

Then, the result is obtained by taking $j = |L|$ in (ii). ☐

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

Lemma 7. *For all $i \geqslant 0$ and for all $\ell \in L$:*

$$\overline{\mathsf{Val}}_\nu^{\leqslant |L|+i}(\ell) \in \mathsf{PossVal}_\nu = [-(|L|-1)\Pi^{\mathrm{tr}} - \Pi^{\mathrm{fin}}, |L|\Pi^{\mathrm{tr}} + \Pi^{\mathrm{fin}}] \cap \mathbf{Z}_{\nu,\boldsymbol{\varphi}}$$

where $\mathsf{PossVal}_\nu$ has cardinality bounded by $|L_f| \times \big((2|L|-1)\Pi^{\mathrm{tr}} + 2\Pi^{\mathrm{fin}} + 1\big)$.

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 $-(|L|-1)\Pi^{\mathrm{tr}} - \Pi^{\mathrm{fin}}$, i.e. $\mathsf{Val}(\ell,\nu) < -(|L|-1)\Pi^{\mathrm{tr}} - \Pi^{\mathrm{fin}}$, then it can secure an arbitrarily small price from that configuration, i.e. $\mathsf{Val}(\ell,\nu) = -\infty$, which contradicts our hypothesis that the value is finite.

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

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

foreach $\ell \in L$ **do**
    if $\ell \in L_f$ **then** $\mathsf{X}(\ell) := \varphi_\ell(\nu)$ **else** $\mathsf{X}(\ell) := +\infty$
repeat
    $\mathsf{X}_{pre} := \mathsf{X}$
    foreach $\ell \in L_{\mathsf{Max}}$ **do** $\mathsf{X}(\ell) := \max_{(\ell,\ell')\in\Delta} \big(\pi(\ell,\ell') + \mathsf{X}_{pre}(\ell')\big)$
    foreach $\ell \in L_{\mathsf{Min}}$ **do** $\mathsf{X}(\ell) := \min_{(\ell,\ell')\in\Delta} \big(\pi(\ell,\ell') + \mathsf{X}_{pre}(\ell')\big)$
    foreach $\ell \in L$ such that $\mathsf{X}(\ell) < -(|L|-1)\Pi^{\mathrm{tr}} - \Pi^{\mathrm{fin}}$ **do** $\mathsf{X}(\ell) := -\infty$
until $\mathsf{X} = \mathsf{X}_{pre}$
return X

Hence, for all $i \geqslant 0$, for all ℓ: $\overline{\mathsf{Val}}_\nu^{\leqslant i}(\ell) \geqslant \mathsf{Val}(\ell,\nu) > -(|L|-1)\Pi^{\mathrm{tr}} - \Pi^{\mathrm{fin}}$. By Lemma 6 and since the sequence is non-increasing, we conclude that, for all $i \geqslant 0$ and for all $\ell \in L$:

$$-(|L|-1)\Pi^{\mathrm{tr}} - \Pi^{\mathrm{fin}} < \overline{\mathsf{Val}}_\nu^{\leqslant |L|+i}(\ell) \leqslant |L|\Pi^{\mathrm{tr}} + \Pi^{\mathrm{fin}}\,.$$

Since all $\overline{\mathsf{Val}}_\nu^{\leqslant |L|+i}(\ell)$ are also in $\mathbf{Z}_{\nu,\boldsymbol{\varphi}}$, we conclude that $\overline{\mathsf{Val}}_\nu^{\leqslant |L|+i}(\ell) \in \mathsf{PossVal}_\nu$ for all $i \geqslant 0$. The upper bound on the size of $\mathsf{PossVal}_\nu$ is established by (2). ☐

This allows us to bound the number of iterations needed for the sequence to stabilise. The worst case is when all locations are assigned a value bounded below by $-(|L|-1)\Pi^{\mathrm{tr}} - \Pi^{\mathrm{fin}}$ from the highest possible values where all vertices are assigned a value bounded above by $|L|\Pi^{\mathrm{tr}} + \Pi^{\mathrm{fin}}$, which is itself reached after $|L|$ steps. Hence:

Corollary 8. *The sequence* $(\overline{\mathsf{Val}}_\nu^{\leqslant i})_{i\geqslant 0}$ *stabilises after a number of steps at most* $|L_f| \times |L| \times \big((2|L|-1)\Pi^{\mathrm{tr}} + 2\Pi^{\mathrm{fin}} + 1\big) + |L|$.

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

Finally, Sections 3.4 of [BGHM16] explain how to compute simultaneously optimal strategies for both players. In our context, this allows us to obtain for every valuation $\nu \in [0, r]$ and location ℓ of an r-SPTG, such that $\mathsf{Val}(\ell,\nu) \notin \{-\infty, +\infty\}$, an optimal FP-strategy for Max, and an optimal switching strategy for Min.

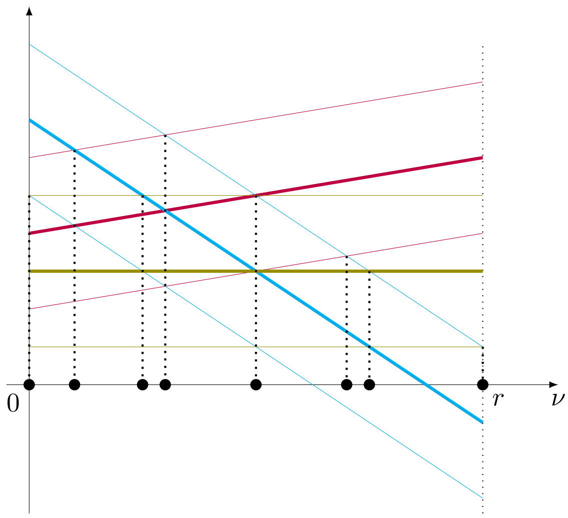

Figure 5: Network of affine functions defined by $\mathsf{F}_{\mathcal{G}}$: functions in bold are final affine functions of $\mathcal{G}$, whereas non-bold ones are their translations with weights $k \in [-(|L|-1)\Pi^{\mathrm{tr}}, |L|\Pi^{\mathrm{tr}}] \cap \mathbf{Z}$. $\mathsf{PossCP}_{\mathcal{G}}$ is the set of abscisses of intersections points, represented by black disks.

4.2. Study of the complete value functions: $\mathcal{G}$ is finitely optimal

Now let us explain how we can reduce the computation of $\mathsf{Val}_{\mathcal{G}}(\ell)\colon \nu \in [0,r] \mapsto \mathsf{Val}(\ell,\nu)$ (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 $\mathsf{CF}_{\{[0,r]\}}$.

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

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

Observe that this set is finite and that its cardinality is $2|L|^2\Pi^{\mathrm{tr}}$, pseudo-polynomial in the size of $\mathcal{G}$. Moreover, as a direct consequence of Corollary 9, this set contains enough information to compute the value of the game in each possible valuation of the clock, in the following sense:

Lemma 10. *For all $\ell \in L$, for all $\nu \in [0,r]$: if $\mathsf{Val}(\ell,\nu)$ is finite, then there is $f \in \mathsf{F}_{\mathcal{G}}$ such that $\mathsf{Val}(\ell,\nu) = f(\nu)$.*

Using the continuity of $\mathsf{Val}_{\mathcal{G}}$ (Theorem 2), this shows that all the cutpoints of $\mathsf{Val}_{\mathcal{G}}$ are intersections of functions from $\mathsf{F}_{\mathcal{G}}$, i.e. belong to the set of *possible cutpoints*

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

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

Proposition 11. *Every r-SPTG $\mathcal{G}$ with only urgent locations is finitely optimal. Moreover, for all locations ℓ, the piecewise affine function $\mathsf{Val}_{\mathcal{G}}(\ell)$ has cutpoints in $\mathsf{PossCP}_{\mathcal{G}}$ of cardinality $4|L_f|^4(\Pi^{\mathrm{tr}})^2$, pseudo-polynomial in the size of $\mathcal{G}$.*

Notice, that this result allows us to compute $\mathsf{Val}(\ell)$ for every $\ell \in L$. First, we compute the set $\mathsf{PossCP}_{\mathcal{G}} = \{y_1, y_2, \dots, y_\ell\}$, which can be done in pseudo-polynomial time in the size of $\mathcal{G}$. Then, for all $1 \leqslant i \leqslant \ell$, we can compute the vectors $\big(\mathsf{Val}(\ell, y_i)\big)_{\ell\in L}$ of values in each location when the clock takes value y_i using Algorithm 1. This provides the value of $\mathsf{Val}(\ell)$ 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 in 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 $\mathsf{Val}_{\mathcal{G}}(\ell)$ is a cost function for all ℓ. This proves item (iii) 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 $\mathsf{PossCP}_{\mathcal{G}}$ in increasing order, and let us assume $\nu_0 = 0$. For all $0 \leqslant i \leqslant k$ let f_i^ℓ be the function from $\mathsf{F}_{\mathcal{G}}$ that defines the piece of $\mathsf{Val}_{\mathcal{G}}(\ell)$ in the interval $[\nu_{i-1}, \nu_i]$ (we have shown above that such an f_i^ℓ always exists). Formally, for all $0 \leqslant i \leqslant k$, $f_i^\ell \in \mathsf{F}_{\mathcal{G}}$ verifies $\mathsf{Val}(\ell, \nu) = f_i^\ell(\nu)$, for all $\nu \in [\nu_{i-1}^\ell, \nu_i^\ell]$. Next, for all $1 \leqslant i \leqslant k$, let μ_i be a value taken in the middle 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 $(\mathsf{Val}_{\mathcal{G}}(\ell, \mu_i))_{\ell\in L}$, and we can extract an optimal memoryless strategy σ^i_{Max} for Max and an optimal switching strategy σ^i_{Min} for Min. Thus we know that, for all $\ell \in L$, playing σ^i_{Min} (respectively, σ^i_{Max}) from (ℓ, μ_i) allows Min (respectively, Max) to ensure a price at most (respectively, at least) $\mathsf{Val}_{\mathcal{G}}(\ell, \mu_i) = f_i^\ell(\mu_i)$. However, it is easy to check that the bound given by $f_i^\ell(\mu_i)$ holds in every valuation, i.e. for all ℓ, for all ν

$$\mathsf{Price}((\ell, \nu), \sigma^i_{\mathsf{Min}}) \leqslant f_i^\ell(\nu) \qquad \text{and} \qquad \mathsf{Price}((\ell, \nu), \sigma^i_{\mathsf{Max}}) \geqslant f_i^\ell(\nu)\,.$$

This holds because:

(i) Min can play σ^i_{Min} from all clock valuations (in $[0, r]$) since we are considering an r-SPTG; and

(ii) Max does not have more possible strategies from an arbitrary valuation $\nu \in [0, r]$ 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^i_{Min} from all configurations (ℓ, ν) with $\nu \in [\nu_{i-1}, \nu_i]$ and secure a price which is at most $f_i^\ell(\nu) = \mathsf{Val}_{\mathcal{G}}(\ell, \nu)$, i.e. σ^i_{Min} is optimal on this interval. By definition of σ^i_{Min}, it is easy to extract from it a fake-optimal NC-strategy (actually, σ^i_{Min} is a switching strategy described by a pair $(\sigma^1_{\mathsf{Min}}, \sigma^2_{\mathsf{Min}})$, and σ^1_{Min} 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 $\mathcal{G} = (L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_f, L_u, \boldsymbol{\varphi}, \Delta, \pi)$ 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 $\mathsf{Val}(\ell, 1)$ for all ℓ assuming all locations are urgent, since time cannot elapse anymore when the clock has valuation 1. This can be done with `solveInstant`. Then, by continuity, $\mathsf{Val}(\ell, 1) = +\infty$ (respectively, $\mathsf{Val}(\ell, 1) = -\infty$) if and only if $\mathsf{Val}(\ell, \nu) = +\infty$ (respectively, $\mathsf{Val}(\ell, \nu) = -\infty$) for all $\nu \in [0, 1]$. We remove from the game all locations with infinite value without changing the values of other locations. Thus, we henceforth assume that $\mathsf{Val}(\ell, \nu) \in \mathbf{R}$ for all (ℓ, ν).

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 *more urgent locations*. Intuitively, a game with more 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, and a valuation $r_0 \in [0, 1]$, we will define a (possibly infinite) sequence of valuations $1 = r_0 > r_1 > \cdots$ and a sequence $\mathcal{G}_{L',r_0}, \mathcal{G}_{L',r_1}, \ldots$ of SPTGs such that

(i) all locations of $\mathcal{G}$ are also present in each $\mathcal{G}_{L',r_i}$, except that the locations of L' are now urgent; and

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

This basic result will be exploited in two directions. First, we prove by induction on the number of urgent locations that all SPTGs are finitely optimal, by re-constructing $\mathsf{Val}_{\mathcal{G}}$ (as well as optimal strategies) as a $\rhd$-concatenation of the value functions of a finite sequence of SPTGs with one more urgent locations. The base case, with only urgent locations, is solved by Proposition 11. This construction suggests a *recursive* algorithm in the spirit of [BLMR06, Rut11] (for non-negative prices). 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 values functions of $\mathcal{G}$ and games $\mathcal{G}_{L',r_i}$.

Let us formalise these constructions. Let $\mathcal{G}$ be an SPTG, let $r \in [0, 1]$ be an endpoint, and let $\boldsymbol{x} = (x_\ell)_{\ell \in L}$ be a vector of rational values. Then, $\mathtt{wait}(\mathcal{G}, r, \boldsymbol{x})$ 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

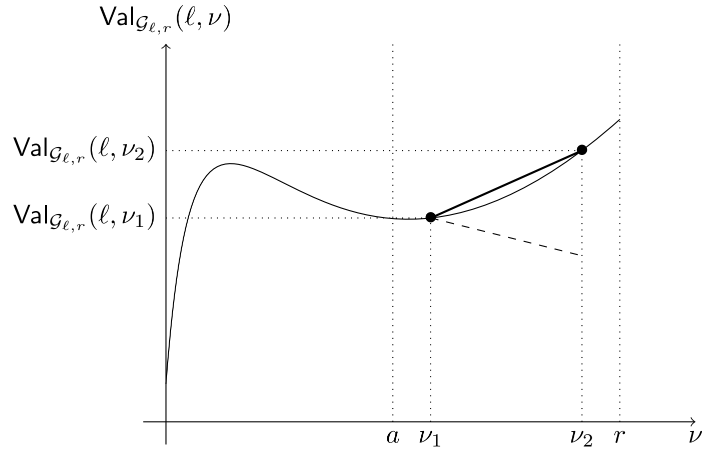

Figure 6: The condition (3) (in the case $L' = \emptyset$ and $\ell \in L_{\mathsf{Min}}$): graphically, it means that the slope between every two points of the plot in $[a, r]$ (represented with a thick line) is greater than or equal to $-\pi(\ell)$ (represented with dashed line).

the game, adding the price x_ℓ to the current price of the play. Formally, $\mathtt{wait}(\mathcal{G}, r, \boldsymbol{x}) = (L_{\mathsf{Min}}, L_{\mathsf{Max}}, L'_f, L_u, \boldsymbol{\varphi}', T', \pi')$ is such that

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

Then, we let $\mathcal{G}_r = \mathtt{wait}\big(\mathcal{G}, r, (\mathsf{Val}_{\mathcal{G}}(\ell, r))_{\ell \in L}\big)$, i.e. the game obtained thanks to $\mathtt{wait}$ by letting $\boldsymbol{x}$ be the value of $\mathcal{G}$ in r. This first transformation does not alter the value of the game, for valuations before r:

Lemma 12. *For all $\nu \in [0, r]$ and locations ℓ,* $\mathsf{Val}_{\mathcal{G}}(\ell, \nu) = \mathsf{Val}_{\mathcal{G}_r}(\ell, \nu)$.

Next, we make locations urgent. For a set $L' \subseteq L \backslash L_u$ 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, introducing urgent locations *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 $\mathsf{Val}_{\mathcal{H}^+}$ as depicted in Figure 6: for each location ℓ of Min, the slopes of the pieces of $\mathsf{Val}_{\mathcal{H}^+}$ contained in $[a, r]$ should be $\leqslant -\pi(\ell)$ (and $\geqslant -\pi(\ell)$ 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 1. Hereafter, we denote the slope of $\mathsf{Val}_{\mathcal{G}}(\ell)$ in-between ν and ν' by $\mathrm{slope}^{\ell}_{\mathcal{G}}(\nu,\nu')$, formally defined by $\mathrm{slope}^{\ell}_{\mathcal{G}}(\nu,\nu') = \frac{\mathsf{Val}_{\mathcal{G}}(\ell,\nu')-\mathsf{Val}_{\mathcal{G}}(\ell,\nu)}{\nu'-\nu}$.

Proposition 13. *Let $0 \leqslant a < r \leqslant 1$, $L' \subseteq L \setminus L_u$ 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 for all $a \leqslant \nu_1 < \nu_2 \leqslant r$*

$$\mathit{slope}^{\ell}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\nu_1,\nu_2) \geqslant -\pi(\ell) \quad (\textit{respectively, } \leqslant -\pi(\ell))\,. \tag{3}$$

Then, for all $\nu \in [a,r]$ and $\ell' \in L$, $\mathsf{Val}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\ell',\nu) = \mathsf{Val}_{\mathcal{G}_{L',r}}(\ell',\nu)$. 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 14. *Let $\mathcal{G}$ be an r-SPTG, ℓ and ℓ' be non-urgent locations of* Min *and* Max*, respectively. Then for all $0 \leqslant \nu < \nu' \leqslant r$:*

$$\mathit{slope}^{\ell}_{\mathcal{G}}(\nu,\nu') \geqslant -\pi(\ell) \qquad \textit{and} \qquad \mathit{slope}^{\ell}_{\mathcal{G}}(\nu,\nu') \leqslant -\pi(\ell')\,.$$

Proof. For the location ℓ, the inequality rewrites in

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

Using the upper definition of the value (thanks to the determinacy result of Theorem 1), it suffices to prove, for all $\varepsilon > 0$, the existence of a strategy σ_{Min} such that for all strategies σ_{Max} of the opponent

$$\mathsf{Price}(\mathsf{Play}((\ell,\nu),\sigma_{\mathsf{Min}},\sigma_{\mathsf{Max}})) \leqslant (\nu'-\nu)\pi(\ell) + \mathsf{Val}_{\mathcal{G}}(\ell,\nu') + \varepsilon\,.$$

First, the definition of the value implies the existence of a strategy σ'_{Min} such that for all strategies σ_{Max}

$$\mathsf{Price}(\mathsf{Play}((\ell,\nu'),\sigma'_{\mathsf{Min}},\sigma_{\mathsf{Max}})) \leqslant \mathsf{Val}_{\mathcal{G}}(\ell,\nu') + \varepsilon\,.$$

Then, σ_{Min} can be obtained by playing from (ℓ,ν), at the first turn, as prescribed by σ'_{Min} but delaying $\nu'-\nu$ time units more (that we are allowed to do since ℓ is non-urgent), and, for other turns, directly like σ'_{Min}. A similar reasoning allows us to obtain the result for ℓ'. □

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 14:

Lemma 15. *For all locations $\ell \in L' \cap L_{\mathsf{Min}}$ (respectively, $\ell \in L' \cap L_{\mathsf{Max}}$), and $\nu \in [0,r]$, $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell,\nu) \leqslant (r-\nu)\pi(\ell) + \mathsf{Val}_{\mathcal{G}}(\ell,r)$ (respectively, $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell,\nu) \geqslant (r-\nu)\pi(\ell) + \mathsf{Val}_{\mathcal{G}}(\ell,r)$).*

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

Proof of Proposition 13. Let σ_{Min} and σ_{Max} be a fake-optimal NC-strategy of Min and an optimal FP-strategy of Max in $\mathcal{G}_{L'\cup\{\ell\},r}$, respectively. 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 $(\ell_0,\nu_0) \xrightarrow{c_0} \cdots \xrightarrow{c_{k-1}} (\ell_k,\nu_k)$, of length $k \geqslant 2$, 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 $\mathsf{int}(\sigma_{\mathsf{Min}})$. For every ℓ_i that is in L_{Min}, and $\nu \in I$, $\sigma_{\mathsf{Min}}(\ell_i,\nu)$ must have a 0 delay, otherwise ν_k would not be in the same interval as ν_0. Thus, the play $(\ell_0,\nu_0) \xrightarrow{c'_0} \cdots \xrightarrow{c'_{k-1}} (\ell_k,\nu_0)$ also conforms with σ_{Min} (with possibly different prices). Furthermore, as all the delays 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, $\pi(\ell_0,\ell_1)+\cdots+\pi(\ell_{k-1},\ell_k) \leqslant -1$, and σ_{Min} is an NC-strategy in $\mathcal{G}_{L',r}$.

We now show the result for $\ell \in L_{\mathsf{Min}}$. The proof for $\ell \in L_{\mathsf{Max}}$ 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 the optimality result of Lemma 4,

$$\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell',\nu) \leqslant \mathsf{fake}^{\sigma_{\mathsf{Min}}}_{\mathcal{G}_{L',r}}(\ell',\nu) = \mathsf{fake}^{\sigma_{\mathsf{Min}}}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\ell',\nu) = \mathsf{Val}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\ell',\nu)\,. \quad (4)$$

To obtain that $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell',\nu) = \mathsf{Val}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\ell',\nu)$, 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 $\mathsf{Price}(\rho) \geqslant \mathsf{Val}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\ell',\nu)$. If ρ has size 1 then ℓ' is a final configuration and $\mathsf{Price}(\rho) = \mathsf{Val}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\ell',\nu) = \varphi'_{\ell'}(\nu)$.

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

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

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

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

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

- 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 $\mathsf{Val}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\ell,\nu'') \leqslant \pi(\ell,\ell'') + \mathsf{Val}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\ell'',\nu'')$. Furthermore, since $\nu'' \in [a,r]$, we can use (3) to obtain

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

Therefore

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

This concludes the induction. As a consequence,

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

for all locations ℓ' and $\nu \in [a, r]$, which finally proves that $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell', \nu) \geqslant \mathsf{Val}_{\mathcal{G}_{L'\cup\{\ell\},r}}(\ell', \nu)$. Fake-optimality of σ_{Min} over $[a, r]$ in $\mathcal{G}_{L'\cup\{\ell\},r}$ is then obtained by (4). ☐

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, with the operator $\mathsf{left}_{L'}\colon (0, 1] \to [0, 1]$ (or simply left, if L' is clear) defined as:

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

By continuity of the value (Theorem 2), this infimum exists and $\mathsf{Val}_{\mathcal{G}}(\ell, \mathsf{left}_{L'}(r)) = \mathsf{Val}_{\mathcal{G}_{L',r}}(\ell, \mathsf{left}_{L'}(r))$. Moreover, $\mathsf{Val}_{\mathcal{G}}(\ell)$ is a cost function on $[\mathsf{left}(r), r]$, since $\mathcal{G}_{L',r}$ is finitely optimal. However, this definition of $\mathsf{left}(r)$ is semantical. Yet, building on the ideas of Proposition 13, we can effectively compute $\mathsf{left}(r)$, given $\mathsf{Val}_{\mathcal{G}_{L',r}}$. We claim that $\mathsf{left}_{L'}(r)$ is the *minimal valuation* such that for all locations $\ell \in L' \cap L_{\mathsf{Min}}$ (respectively, $\ell \in L' \cap L_{\mathsf{Max}}$), the slopes of the affine sections of the cost function $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell)$ on $[\mathsf{left}(r), r]$ are at least (at most) $-\pi(\ell)$. Hence, $\mathsf{left}(r)$ can be obtained (see Figure 7) by inspecting iteratively, for all ℓ of Min (respectively, Max), the slopes of $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell)$ by decreasing valuations 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 $\mathsf{Val}_{\mathcal{G}_{L',r}}$ has finitely many pieces, by hypothesis. Moreover, this guarantees that $\mathsf{left}(r) < r$, as shown in the following lemma.

Lemma 16. *Let $\mathcal{G}$ be an SPTG, $L' \subseteq L \setminus L_u$, and $r \in (0, 1]$, such that $\mathcal{G}_{L'',r}$ is finitely optimal for all $L'' \subseteq L'$. Then, $\mathsf{left}_{L'}(r)$ is the minimal valuation such that for all locations $\ell \in L' \cap L_{\mathsf{Min}}$ (respectively, $\ell \in L' \cap L_{\mathsf{Max}}$), the slopes of the affine sections of the cost function $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell)$ on $[\mathsf{left}(r), r]$ are at least (respectively, at most) $-\pi(\ell)$. Moreover, $\mathsf{left}(r) < r$.*

Proof. Since $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell) = \mathsf{Val}_{\mathcal{G}}(\ell)$ on $[\mathsf{left}(r), r]$, and as ℓ is non-urgent in $\mathcal{G}$, Lemma 14 states that all the slopes of $\mathsf{Val}_{\mathcal{G}}(\ell)$ are at least (respectively, at most) $-\pi(\ell)$ on $[\mathsf{left}(r), r]$.

We now show the minimality property by contradiction. Therefore, let $r' < \mathsf{left}(r)$ such that all cost functions $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell)$ are affine on $[r', \mathsf{left}(r)]$, and assume that for all $\ell \in L' \cap L_{\mathsf{Min}}$ (respectively, $\ell \in L' \cap L_{\mathsf{Max}}$), the slopes of $\mathsf{Val}_{\mathcal{G}_{L',r}}(\ell)$ on $[r', \mathsf{left}(r)]$ are at least (respectively, at most) $-\pi(\ell)$. Hence, this property holds on $[r', r]$. Then, by applying Proposition 13 $|L'|$ times (here, we use the finite optimality of the games $\mathcal{G}_{L'',r}$ with $L'' \subseteq L'$), we have that for all $\nu \in [r', r]$ $\mathsf{Val}_{\mathcal{G}_r}(\ell, \nu) = \mathsf{Val}_{\mathcal{G}_{L',r}}(\ell, \nu)$. Using Lemma 12, we also know that for all $\nu \leqslant r$,

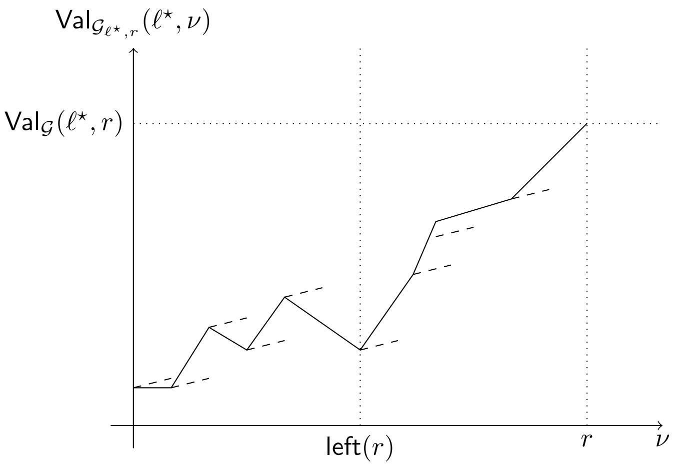

Figure 7: In this example $L' = \{\ell^\star\}$ and $\ell^\star \in L_{\mathsf{Min}}$. $\mathsf{left}(r)$ is the leftmost point such that all slopes on its right are smaller than or equal to $-\pi(\ell^\star)$ in the graph of $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star,\nu)$. Dashed lines have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

and ℓ, $\mathsf{Val}_{\mathcal{G}_r}(\ell,\nu) = \mathsf{Val}_{\mathcal{G}}(\ell,\nu)$. Thus, $\mathsf{Val}_{\mathcal{G}_{r,L'}}(\ell,\nu) = \mathsf{Val}_{\mathcal{G}}(\ell,\nu)$. As $r' < \mathsf{left}(r)$, this contradicts the definition of $\mathsf{left}_{L'}(r)$.

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

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

From Lemma 15, we also know that

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

Both equations combined imply

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

which is not possible if $\mathsf{left}(r) = r$. □

Thus, one can reconstruct $\mathsf{Val}_{\mathcal{G}}$ on $[\inf_i r_i, r_0]$ from the value functions of the (potentially infinite) sequence of games $\mathcal{G}_{L',r_0}$, $\mathcal{G}_{L',r_1}, \ldots$ where $r_{i+1} = \mathsf{left}(r_i)$ for all i such that $r_i > 0$, for all possible choices of non-urgent locations L'. Next, we will define two different ways of choosing L': the former to prove finite optimality of all SPTGs, the latter to obtain an 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 $L' = \{\ell^\star\}$ where $\ell^\star$ is a non-urgent location of *minimum price-rate* (i.e. for all $\ell \in L$, $\pi(\ell^\star) \leqslant \pi(\ell)$). Given $r_0 \in [0,1]$, we let $r_0 > r_1 > \cdots$ be the decreasing sequence of valuations such that $r_i = \mathsf{left}_{\ell^\star}(r_{i-1})$ for all $i > 0$. As explained before, we will build $\mathsf{Val}_{\mathcal{G}}$ on $[\inf_i r_i, r_0]$ from the value functions of games $\mathcal{G}_{\ell^\star, r_i}$. Assuming finite optimality of those games, this will prove that $\mathcal{G}$ is finitely optimal *under the condition* that $r_0 > r_1 > \cdots$ eventually stops, i.e. $r_i = 0$ for some i. Lemma 18 will prove this property. First, we relate the optimal value functions with the final cost functions.

Lemma 17. *Assume that $\mathcal{G}_{\ell^\star,r}$ is finitely optimal. If $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star)$ is affine on a non-singleton interval $I \subseteq [0,r]$ with a slope greater*[6] *than $-\pi(\ell^\star)$, then there exists $f \in \mathsf{F}_{\mathcal{G}}$ (see definition in page 19) such that for all $\nu \in I$, $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star, \nu) = f(\nu)$.*

Proof. Let σ_{Min} and σ_{Max} be some fake-optimal NC-strategy and optimal FP-strategy in $\mathcal{G}_{\ell^\star,r}$. 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} and of σ_{Max}. Let σ'_{Min} the optimal switching strategy obtained from σ_{Min} in Lemma 4: notice that both strategies have the same intervals.

Let $\nu \in I'$. Since $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star, \nu) < +\infty$, the play $\mathsf{Play}((\ell^\star, \nu), \sigma'_{\mathsf{Min}}, \sigma_{\mathsf{Max}})$ necessarily reaches a final location and has price $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star, \nu)$. Let $(\ell_0, \nu_0) \xrightarrow{c_0} \cdots (\ell_k, \nu_k)$ be its prefix until the first final location ℓ_k (the prefix used to compute the price of the play). We also let $\nu' \in I'$ be a valuation such that $\nu < \nu'$.

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 $\ell_j \in L_{\mathsf{Min}}$, let $(t, \delta) = \sigma'_{\mathsf{Min}}(\ell_j, \nu)$ and $(t', \delta') = \sigma'_{\mathsf{Min}}(\ell_j, \nu')$. Similarly, if $\ell_j \in L_{\mathsf{Max}}$, we let $(t, \delta) = \sigma_{\mathsf{Max}}(\ell_j, \nu)$ and $(t', \delta') = \sigma_{\mathsf{Max}}(\ell_j, \nu')$. As I' is contained in an interval of σ'_{Min} and σ_{Max}, we have $\delta = \delta'$ 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 $\mathsf{Play}((\ell^\star, \nu'), \sigma'_{\mathsf{Min}}, \sigma_{\mathsf{Max}})$: notice moreover that, as before, this prefix has price $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star, \nu')$. In particular,

$$\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star, \nu') = \mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star, \nu) - (\nu' - \nu)\pi(\ell_{i-1}) \leqslant \mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star, \nu) - (\nu' - \nu)\pi(\ell^\star)$$

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

Again by contradiction, assume now that $\ell_k = \ell^f$ for some $\ell \in L \setminus L_u$. By the same reasoning as before, we then would have $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star, \nu') = \mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star, \nu) - (\nu' - \nu)\pi(\ell)$, which again contradicts the hypothesis.

Therefore, $\ell_k \in L_f$. 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}: since it is an NC-strategy,

[6] For this result, the order does not depend on the owner of the location, but on the fact that $\ell^\star$ has minimal price amongst locations of $\mathcal{G}$.

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 he wants, and thus $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star,\nu) = -\infty$. 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. Thus, the sum of the discrete weights $w = \pi(\ell_0,\ell_1) + \cdots + \pi(\ell_{k-1},\ell_k)$ belongs to the set $[-(|L|-1)\Pi^{\mathrm{tr}}, |L|\Pi^{\mathrm{tr}}] \cap \mathbf{Z}$, and we have $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star,\nu) = w + \varphi_{\ell_k}(\nu)$. Notice that the previous developments also show that for all $\nu' \in I'$ (here, $\nu < \nu'$ is not needed), $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star,\nu') = w + \varphi_{\ell_k}(\nu')$, with the same location ℓ_k, and weight k. Since this equality holds on $I' \subseteq I$ which is not a singleton, and $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r}}(\ell^\star)$ is affine on I, it holds everywhere on I. This shows the result since $w + \varphi_{\ell_k} \in \mathsf{F}_{\mathcal{G}}$. ☐

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 $\ell^\star$ has a strictly better strategy in configuration $(\ell^\star, r_{i+1})$ than waiting until r_i in location $\ell^\star$.

Lemma 18. *If $\mathcal{G}_{\ell^\star,r_i}$ is finitely optimal for all $i \geqslant 0$, then*

(i) *if $\ell^\star \in L_{\mathsf{Min}}$ (respectively, L_{Max}), $\mathsf{Val}_{\mathcal{G}}(\ell^\star, r_{i+1}) < \mathsf{Val}_{\mathcal{G}}(\ell^\star, r_i) + (r_i - r_{i+1})\pi(\ell^\star)$ (respectively, $\mathsf{Val}_{\mathcal{G}}(\ell^\star, r_{i+1}) > \mathsf{Val}_{\mathcal{G}}(\ell^\star, r_i) + (r_i - r_{i+1})\pi(\ell^\star)$), for all i; and*

(ii) *there is $i \leqslant |\mathsf{F}_{\mathcal{G}}|^2 + 2$ such that $r_i = 0$.*

Proof. For the first item, we assume $\ell^\star \in L_{\mathsf{Min}}$, since the proof of the other case only differ with respect to the sense of the inequalities. From Lemma 16, we know that in $\mathcal{G}_{\ell^\star,r_i}$ there exists $r' < r_{i+1}$ such that $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r_i}}(\ell^\star)$ is affine of $[r', r_{i+1}]$ and its slope is smaller that $-\pi(\ell^\star)$, i.e. $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r_i}}(\ell^\star, r_{i+1}) < \mathsf{Val}_{\mathcal{G}_{\ell^\star,r_i}}(r') - (r_{i+1} - r')\pi(\ell^\star)$. Lemma 15 also ensures that $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r_i}}(\ell^\star, r') \leqslant \mathsf{Val}_{\mathcal{G}}(\ell^\star, r_i) + (r_i - r')\pi(\ell^\star)$. Combining both inequalities allows us to conclude.

We now turn to the proof of the second item, showing the stationarity of sequence $(r_i)_{i\geqslant 0}$. We consider first the case where $\underline{\ell^\star \in L_{\mathsf{Max}}}$. Let $i > 0$ such that $r_i \neq 0$ (if there exist no such i then $r_1 = 0$). Recall from Lemma 16 that there exists $r'_i < r_i$ such that $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r_{i-1}}}(\ell^\star)$ is affine on $[r'_i, r_i]$, of slope greater than $-\pi(\ell^\star)$. In particular,

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

Lemma 17 states that on $[r'_i, r_i]$, $\mathsf{Val}_{\mathcal{G}_{\ell^\star,r_{i-1}}}(\ell^\star)$ is equal to some $f_i \in \mathsf{F}_{\mathcal{G}}$. As f_i is an affine function, $f_i(r_i) = \mathsf{Val}_{\mathcal{G}_{\ell^\star,r_{i-1}}}(\ell^\star, r_i)$, and $f_i(r'_i) = \mathsf{Val}_{\mathcal{G}_{\ell^\star,r_{i-1}}}(\ell^\star, r'_i)$. Thus, for all ν,

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

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

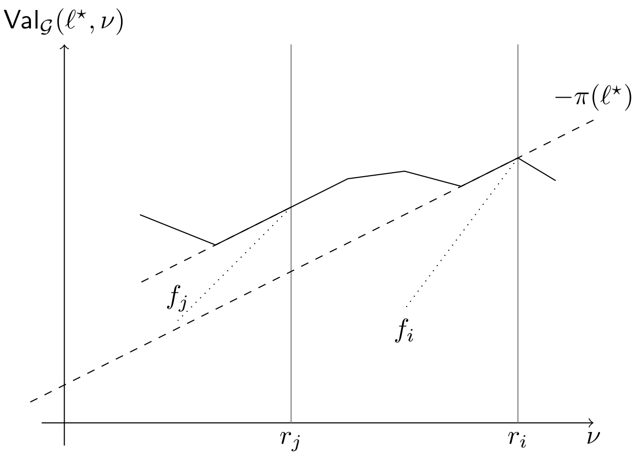

Figure 8: The case $\ell^\star \in L_{\mathsf{Max}}$: a geometric proof of $f_i \neq f_j$. The dotted lines represents f_i and f_j, the dashed lines have slope $-\pi(\ell^\star)$, and the plain line depicts $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$. Because the slope of f_i is strictly smaller than $-\pi(\ell^\star)$,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

Consider then $j > i$ such that $r_j \neq 0$. We claim that $f_j \neq f_i$. Indeed, we have $\mathsf{Val}_{\mathcal{G}}(\ell^\star, r_j) = f_j(r_j)$. As, in $\mathcal{G}$, $\ell^\star$ is a non-urgent location, Lemma 14 ensures that

$$\mathsf{Val}_{\mathcal{G}}(\ell^\star, r_j) \geqslant \mathsf{Val}_{\mathcal{G}}(\ell^\star, r_i) + \pi(\ell^\star)(r_i - r_j)\,.$$

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

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

We continue with the case where $\underline{\ell^\star \in L_{\mathsf{Min}}}$. Let $r_\infty = \inf\{r_i \mid i \geqslant 0\}$. In this case, we look at the affine parts of $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$ with a slope greater than $-\pi(\ell^\star)$, 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.

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 $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$ over $(r_\infty, 1]$. Recall that it means that $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$ is affine on $[a, b]$, and if we let a' be the greatest cutpoint smaller than a, and b' the smallest cutpoint greater than b, the slopes of $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$ 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 $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$ on $[a, b]$ and simply call *cutpoint* a cutpoint of $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$.

To every segment $[a, b]$ with a slope greater than $-\pi(\ell^\star)$, we associate a function $f_{[a,b]} \in \mathsf{F}_{\mathcal{G}}$ 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 17 ensures that there exists $f_{[a,b]} \in \mathsf{F}_{\mathcal{G}}$ such that for all $\nu \in [a', b']$, $\mathsf{Val}_{\mathcal{G}}(\ell^\star, \nu) = f_{[a,b]}(\nu)$.

Consider now two disjoint segments $[a, b]$ and $[c, d]$ with a slope strictly greater than $-\pi(\ell^\star)$, and assume that $f_{[a,b]} = f_{[c,d]}$ (in particular both segments

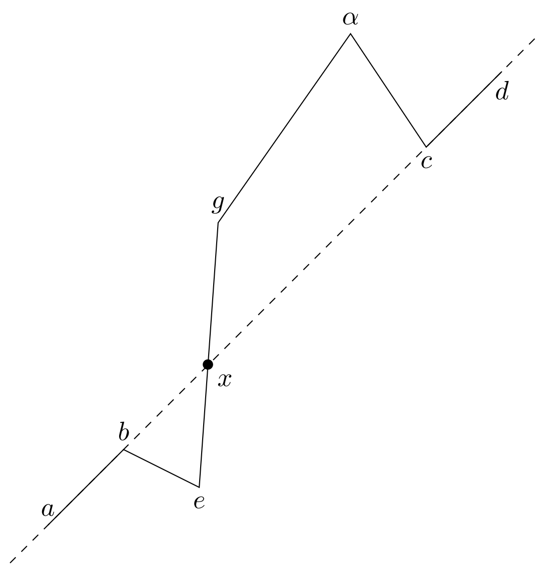

Figure 9: In order for the segments $[a, b]$ and $[c, d]$ to be aligned, there must exist a segment with a biggest slope crossing $f_{[a,b]}$ (represented by a dashed line) between b and c.

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 $[b, c]$, in a point of abscisse x, i.e. $x \in [b, c]$ verifies $f_{[e,g]}(x) = f_{[a,b]}(x)$ (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 $\nu \in [\alpha, c)$, $\mathsf{Val}_{\mathcal{G}}(\ell^\star, \nu) > f_{[c,d]}(\nu)$. Let x be the greatest point in $[b, \alpha]$ such that $\mathsf{Val}_{\mathcal{G}}(\ell^\star, x) = f_{[c,d]}(x)$. We know that it exists since $\mathsf{Val}_{\mathcal{G}}(\ell^\star, b) = f_{[c,d]}(b)$, and $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$ is continuous. Observe that $\mathsf{Val}_{\mathcal{G}}(\ell^\star, \nu) > f_{[c,d]}(\nu)$, for all $x < \nu < c$. Finally, let g be the smallest cutpoint of $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$ strictly greater than x, and e the greatest cutpoint of $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$ smaller than or equal to x. By construction $[e, g]$ is a segment that contains x. The slope of the segment $[e, g]$ is $s_{[e,g]} = \frac{\mathsf{Val}_{\mathcal{G}}(\ell^\star,g)-\mathsf{Val}_{\mathcal{G}}(\ell^\star,x)}{g-x}$, 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 $\mathsf{Val}_{\mathcal{G}}(\ell^\star, x) = f_{[c,d]}(x)$, and that $\mathsf{Val}_{\mathcal{G}}(\ell^\star, g) > f_{[c,d]}(g)$ since $g \in (x, c)$, we obtain that $s_{[e,g]} > s_{[c,d]}$. Finally, since $\mathsf{Val}_{\mathcal{G}}(\ell^\star, x) = f_{[c,d]}(x) = f_{[e,g]}(x)$, it is indeed the abscisse of 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 $f \in \mathsf{F}_{\mathcal{G}}$, there are less than $|\mathsf{F}_{\mathcal{G}}|$ intersection points between f and the other functions of $\mathsf{F}_{\mathcal{G}}$ (at most one for each pair (f, f')). If f has a slope greater than $-\pi(\ell^\star)$, thanks to the previous paragraph, we know that there are at most $|\mathsf{F}_{\mathcal{G}}|$ segments $[a, b]$ such that $f_{[a,b]} = f$. Summing over all possible functions f, there are at most $|\mathsf{F}_{\mathcal{G}}|^2$ segments with a slope greater than $-\pi(\ell^\star)$.

Now, we link those segments with the valuations r_i's, for $i > 0$. By item

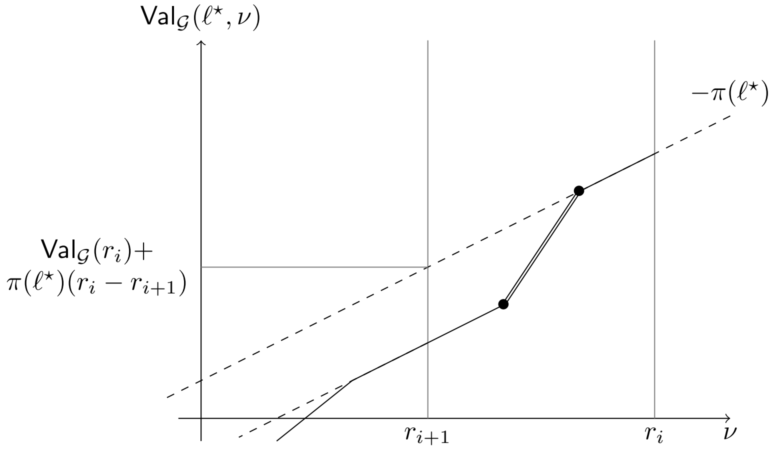

Figure 10: The case $\ell^\star \in L_{\mathsf{Min}}$: as the value at r_{i+1} is strictly below $\mathsf{Val}_{\mathcal{G}}(r_i)+\pi(\ell^\star)(r_i-r_{i+1})$, as the slope on the left of r_i and of r_{i+1} is $-\pi(\ell^\star)$, there must exist a segment (represented with a double line) with slope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

(i), thanks to the finite-optimality of $\mathcal{G}_{\ell^\star, r_i}$, $\mathsf{Val}_{\mathcal{G}}(\ell^\star, r_{i+1}) < (r_i - r_{i+1})\pi(\ell^\star) + \mathsf{Val}_{\mathcal{G}}(\ell^\star, r_i)$. Furthermore, Proposition 20 states that the slope of the segment directly on the left of r_i is equal to $-\pi(\ell^\star)$. With the previous inequality in mind, this cannot be the case if $\mathsf{Val}_{\mathcal{G}}(\ell^\star)$ is affine over the whole interval $[r_{i+1}, r_i]$. Thus, there exists a segment $[a, b]$ of slope strictly greater than $-\pi(\ell^\star)$ such that $b \in [r_{i+1}, r_i]$. As we also know that the slope left to r_{i+1} is $-\pi(\ell^\star)$, 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 $|\mathsf{F}_{\mathcal{G}}|^2$, we know that the sequence r_i is stationary in at most $|\mathsf{F}_{\mathcal{G}}|^2 + 1$ steps, i.e. that there exists $i \leqslant |\mathsf{F}_{\mathcal{G}}|^2 + 1$ such that $r_i = 0$. $\square$

By iterating this construction, we make all locations urgent iteratively, and obtain:

Theorem 19. *Every SPTG $\mathcal{G}$ is finitely optimal and for all locations ℓ, $\mathsf{Val}_{\mathcal{G}}(\ell)$ has at most $O\left((\Pi^{\mathrm{tr}}|L|^2)^{2|L|+2}\right)$ cutpoints.*

Proof. As announced, we show by induction on $n \geqslant 0$ that every r-SPTG $\mathcal{G}$ with n non-urgent locations is finitely optimal, and that the number of cutpoints of $\mathsf{Val}_{\mathcal{G}}(\ell)$ is at most $O\left((\Pi^{\mathrm{tr}}(|L_f| + n^2))^{2n+2}\right)$, which suffices to show the above bound, since $|L_f| + n^2 \leqslant |L|^2$.

The base case $n = 0$ is given by Proposition 11. Now, assume that $\mathcal{G}$ has at least one non-urgent location, and consider $\ell^\star$ one with minimum price-rate. By induction hypothesis, all r'-SPTGs $\mathcal{G}_{\ell^\star, r'}$ are finitely optimal for all $r' \in [0, r]$. Let $r_0 > r_1 > \cdots$ be the decreasing sequence defined by $r_0 = r$ and $r_i = \mathsf{left}_{\ell^\star}(r_{i-1})$ for all $i \geqslant 1$. By Lemma 18, there exists $j \leqslant |\mathsf{F}_{\mathcal{G}}|^2 + 2$ such that $r_j = 0$. Moreover, for all $0 < i \leqslant j$, $\mathsf{Val}_{\mathcal{G}} = \mathsf{Val}_{\mathcal{G}_{\ell^\star, r_{i-1}}}$ on $[r_i, r_{i-1}]$ by definition of $r_i = \mathsf{left}_{\ell^\star}(r_{i-1})$, so that $\mathsf{Val}_{\mathcal{G}}(\ell)$ is a cost function on this interval, for all ℓ, and the number of cutpoints on this interval is bounded by $O\left((\Pi^{\mathrm{tr}}(|L_f| + (n-1)^2 + n))^{2(n-1)+2}\right) = O\left((\Pi^{\mathrm{tr}}(|L_f| + n^2))^{2(n-1)+2}\right)$ by induction hypothesis (notice that maximal transition prices are the same in $\mathcal{G}$ and

$\mathcal{G}_{\ell^\star,r_{i-1}}$, but that we add n more final locations in $\mathcal{G}_{\ell^\star,r_{i-1}}$). Adding the cutpoint 1, summing over i from 0 to $j \leqslant |\mathsf{F}_\mathcal{G}|^2+2$, and observing that $|\mathsf{F}_\mathcal{G}| \leqslant 2\Pi^{\mathrm{tr}}|L_f|$, we bound the number of cutpoints of $\mathsf{Val}_\mathcal{G}(\ell)$ by $O\left((\Pi^{\mathrm{tr}}(|L_f|+n^2))^{2n+2}\right)$. Finally, we can reconstruct fake-optimal and optimal strategies in $\mathcal{G}$ from the fake-optimal and optimal strategies of $\mathcal{G}_{\ell^\star,r_i}$. $\square$

6. Algorithms to compute the value function

The finite optimality of SPTGs allows us to compute the value functions. The proof of Theorem 19 suggests a *recursive* algorithm to do so: from an SPTG $\mathcal{G}$ with minimal non-urgent location $\ell^\star$, solve recursively $\mathcal{G}_{\ell^\star,1}$, $\mathcal{G}_{\ell^\star,\mathsf{left}(1)}$, $\mathcal{G}_{\ell^\star,\mathsf{left}(\mathsf{left}(1))}$, etc. handling the base case where all locations are urgent with Algorithm 1. While our results above show that this is correct and terminates, we propose instead to solve—without the need for recursion—the sequence of games $\mathcal{G}_{L\setminus L_u,1}$, $\mathcal{G}_{L\setminus L_u,\mathsf{left}(1)},\ldots$ i.e. *making all locations urgent at once.* Again, the arguments given above prove that this scheme is *correct*, but the key argument of Lemma 18 that ensures *termination* cannot be applied in this case. Instead, we rely on the following result, stating, that there will be at least one cutpoint of $\mathsf{Val}_\mathcal{G}$ in each interval $[\mathsf{left}(r), r]$. Observe that this lemma relies on the fact that $\mathcal{G}$ is finitely optimal, hence the need to first prove this fact independently with the sequence $\mathcal{G}_{\ell^\star,1}$, $\mathcal{G}_{\ell^\star,\mathsf{left}(1)}$, $\mathcal{G}_{\ell^\star,\mathsf{left}(\mathsf{left}(1))},\ldots$ Termination then follows from the fact that $\mathsf{Val}_\mathcal{G}$ has finitely many cutpoints by finite optimality.

Proposition 20. *Let $r_0 \in (0,1]$ such that $\mathcal{G}_{L',r_0}$ is finitely optimal. Suppose that $r_1 = \mathsf{left}_{L'}(r_0) > 0$, and let $r_2 = \mathsf{left}_{L'}(r_1)$. There exists $r' \in [r_2, r_1)$ and $\ell \in L'$ such that*

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

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

Proof. We denote by r' the smallest valuation (smaller than r_1) such that for all locations ℓ, $\mathsf{Val}_\mathcal{G}(\ell)$ is affine over $[r', r_1]$. Then, the proof goes by contradiction: using Lemma 16, we assume that for all $\ell \in L'\cap L_{\mathsf{Min}}$ (respectively, $\ell \in L'\cap L_{\mathsf{Max}}$)

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

Let σ^0_{Min} and σ^0_{Max} (respectively, σ^1_{Min} and σ^1_{Max}) be a fake-optimal NC-strategy and an optimal FP-strategy in $\mathcal{G}_{L',r_0}$ (respectively, $\mathcal{G}_{L',r_1}$). Let $r'' = \max(\mathsf{pts}(\sigma^1_{\mathsf{Min}}) \cup \mathsf{pts}(\sigma^1_{\mathsf{Max}})) \cap [r', r_1)$, so that strategies σ^1_{Min} and σ^1_{Max} have the *same behaviour* on all valuations 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 valuation greater than or equal to r_1 and then play the same transition.

Observe preliminarily that for all $\ell \in L' \cap L_{\mathsf{Min}}$ (respectively, $\ell \in L' \cap L_{\mathsf{Max}}$), if on the interval (r'', r_1), σ^1_{Min} (respectively, σ^1_{Max}) goes to ℓ^f then the slope on $[r'', r_1]$ (and thus on $[r', r_1]$) is $-\pi(\ell)$. Thus for such a location ℓ, we know that $(i) \wedge \neg(ii)$ holds for ℓ (by letting r' be r'').

For other locations ℓ, 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 valuations $\nu \in (r'', r_1)$

$$\mathsf{fake}^{\sigma_{\mathsf{Min}}}_{\mathcal{G}_{L',r_0}}(\ell,\nu) \leqslant \mathsf{Val}_{\mathcal{G}}(\ell,\nu) \leqslant \mathsf{Price}_{\mathcal{G}_{L',r_0}}((\ell,\nu),\sigma_{\mathsf{Max}}) \,. \tag{5}$$

As a consequence, with Lemma 4 (over game $\mathcal{G}_{L',r_0}$), one would have that $\mathsf{Val}_{\mathcal{G}_{L',r_0}}(\ell,\nu) = \mathsf{Val}_{\mathcal{G}}(\ell,\nu)$, which will raise a contradiction with the definition of r_1 as $\mathsf{left}_{L'}(r_0) < r_0$ (by Lemma 16), 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^1_{Min} over $[0, r_1]$, and σ^0_{Min} over $[r_1, r_0]$: a special care has to be spent in case σ^1_{Min} performs a jump to a location ℓ^f, since then, in σ_{Min}, we rather glue this move with the decision of strategy σ^0_{Min} in (ℓ, r_1). Formally, let (ℓ,ν) be a configuration of $\mathcal{G}_{L',r_0}$ with $\ell \in L_{\mathsf{Min}}$. We construct $\sigma_{\mathsf{Min}}(\ell,\nu)$ as follows:

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

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 $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leqslant \mathsf{Val}_{\mathcal{G}}(\ell,\nu)$. Note that ρ either only contains valuations 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 run that satisfies the above restriction.

- If $\nu \in [r_1, r_0]$, then ρ conforms with σ^0_{Min}, thus, as σ^0_{Min} is fake-optimal, $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leqslant \mathsf{Val}_{\mathcal{G}_{L',r_0}}(\ell,\nu) = \mathsf{Val}_{\mathcal{G}}(\ell,\nu)$ (the last inequality comes from the definition of $r_1 = \mathsf{left}_{L'}(r_0)$). 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 $\ell \in L' \cap L_{\mathsf{Min}}$, 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^1_{Min} on (r'', r_1) consists in going to ℓ^f, thus, as observed above, $(i) \wedge \neg(ii)$ holds for ℓ. Therefore, $\mathsf{Val}_{\mathcal{G}}(\ell,\nu) = \mathsf{Val}_{\mathcal{G}}(\ell,r_0) + \pi(\ell)(r_0 - \nu) = \varphi_{\ell_f}(\nu) = \mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho)$.
  - If $\ell \in L_{\mathsf{Min}} \setminus L'$, by construction, it must be the case that $\sigma_{\mathsf{Min}}(\ell,\nu) = \big(r_1 - \nu + t', (\ell,\ell^f)\big)$ where $\big(t,(\ell,\ell^f)\big) = \sigma^1_{\mathsf{Min}}(\ell,\nu)$ and $\big(t',(\ell,\ell^f)\big) = \sigma^0_{\mathsf{Min}}(\ell,r_1)$. Thus, $\nu' = r_1 + t'$. In particular, observe that $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho) = (r_1-\nu)\pi(\ell) + \mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho')$ where $\rho' = (\ell,r_1) \xrightarrow{c'} (\ell^f,\nu')$. As ρ' conforms with σ^0_{Min} which is fake-optimal in $\mathcal{G}_{L',r_0}$, $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho') \leqslant \mathsf{Val}_{\mathcal{G}_{L',r_0}}(\ell,r_1) = \mathsf{Val}_{\mathcal{G}}(\ell,r_1)$ (since $r_1 = \mathsf{left}(r_0)$). Thus $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leqslant (r_1-\nu)\pi(\ell) + \mathsf{Val}_{\mathcal{G}}(\ell,r_1) = \mathsf{Price}_{\mathcal{G}_{L',r_1}}(\rho'')$ where $\rho'' = (\ell,\nu) \xrightarrow{c''} (\ell^f, \nu+t)$ conforms with σ^1_{Min} which is fake-optimal in $\mathcal{G}_{L',r_1}$. Therefore, $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leqslant \mathsf{Val}_{\mathcal{G}_{L',r_1}}(\ell,\nu) = \mathsf{Val}_{\mathcal{G}}(\ell,\nu)$ (since $r_1 = \mathsf{left}(r_0)$).

– If $\ell \in L_{\mathsf{Max}}$ then $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho) = (\nu' - \nu)\pi(\ell) + \varphi_{\ell_f}(\nu') = (\nu' - \nu)\pi(\ell) + (r_0 - \nu')\pi(\ell) + \mathsf{Val}_{\mathcal{G}}(\ell, r_0) = (r_0 - \nu)\pi(\ell) + \mathsf{Val}_{\mathcal{G}}(\ell, r_0)$. By Lemma 14, since $\ell \in L_{\mathsf{Max}} \setminus L_u$ (ℓ is not urgent in $\mathcal{G}$ since ℓ^f exists), $\mathsf{Val}_{\mathcal{G}}(\ell, r_1) \geqslant (r_0 - r_1)\pi(\ell) + \mathsf{Val}_{\mathcal{G}}(\ell, r_0)$. 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^1_{Min} and

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

Thus, as σ^1_{Min} is fake-optimal in $\mathcal{G}_{L',r_1}$, $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leqslant \mathsf{Price}_{\mathcal{G}_{L',r_1}}(\rho') \leqslant \mathsf{Val}_{\mathcal{G}_{L',r_1}}(\ell, \nu) = \mathsf{Val}_{\mathcal{G}}(\ell, \nu)$.

- 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 $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho') \leqslant \mathsf{Val}_{\mathcal{G}}(\ell', \nu')$.

  – If $\nu' \leqslant r_1$, let ρ'' be the play of $\mathcal{G}_{L',r_1}$ starting in (ℓ', ν') that conforms with σ^1_{Min} and σ^1_{Max}. If ρ'' does not reach a final location, since σ^1_{Min} is an NC-strategy, the prices of its prefixes tend to $-\infty$. By considering the strategy σ'_{Min} of Lemma 4, we would obtain a run conforming with σ^1_{Max} of price smaller than $\mathsf{Val}_{\mathcal{G}_{L',r_1}}(\ell', \nu')$ which would contradict the optimality of σ^1_{Max}. Hence, ρ'' reaches the target. Moreover, since σ^1_{Max} is optimal and σ^1_{Min} is fake-optimal, we finally know that $\mathsf{Price}_{\mathcal{G}_{L',r_1}}(\rho'') = \mathsf{Val}_{\mathcal{G}_{L',r_1}}(\ell', \nu') = \mathsf{Val}_{\mathcal{G}}(\ell', \nu')$ (since $\nu' \in [\mathsf{left}(r_1), r_1]$). Therefore,

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

  Since the play $(\ell, \nu) \xrightarrow{c'} \rho''$ conforms with σ^1_{Min}, we finally have $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leqslant \mathsf{Price}((\ell, \nu) \xrightarrow{c'} \rho'') \leqslant \mathsf{Val}_{\mathcal{G}_{L',r_1}}(\ell, \nu) = \mathsf{Val}_{\mathcal{G}}(\ell, \nu)$.

  – If $\nu' > r_1$ and $\ell \in L_{\mathsf{Max}}$, 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

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

  Since ρ^0 conforms with σ^0_{Min}, fake-optimal, and reaches a final location, $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho^0) \leqslant \mathsf{Val}_{\mathcal{G}_{L',r_0}}(\ell, r_1) = \mathsf{Val}_{\mathcal{G}}(\ell, r_1)$ (since $r_1 =$

Algorithm 2: `solve`($\mathcal{G}$)

Input: SPTG $\mathcal{G} = (L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_f, L_u, \boldsymbol{\varphi}, \Delta, \pi)$

1 $\boldsymbol{f} = (f_\ell)_{\ell \in L} :=$ `solveInstant`$(\mathcal{G}, 1)$ `/*` $f_\ell \colon \{1\} \to \overline{\mathbf{R}}$ `*/`

2 $r := 1$

3 **while** $0 < r$ **do** `/* Invariant:` $f_\ell \colon [r, 1] \to \overline{\mathbf{R}}$ `*/`

4 $\quad$ $\mathcal{G}' :=$ `wait`$(\mathcal{G}, r, \boldsymbol{f}(r))$ `/*` r-SPTG $\mathcal{G}' = (L_{\mathsf{Min}}, L_{\mathsf{Max}}, L'_f, L'_u, \boldsymbol{\varphi}', T', \pi')$ `*/`

5 $\quad$ $L'_u := L'_u \cup L$ `/* every location is made urgent */`

6 $\quad$ $b := r$

7 $\quad$ **repeat** `/* Invariant:`$f_\ell \colon [b, 1] \to \overline{\mathbf{R}}$ `*/`

8 $\quad\quad$ $a := \max(\mathsf{PossCP}_{\mathcal{G}'} \cap [0, b))$

9 $\quad\quad$ $\boldsymbol{x} = (x_\ell)_{\ell \in L} :=$ `solveInstant`$(\mathcal{G}', a)$ `/*` $x_\ell = \mathsf{Val}_{\mathcal{G}'}(\ell, a)$ `*/`

10 $\quad\quad$ **if** $\forall \ell \in L_{\mathsf{Min}}\ \frac{f_\ell(b) - x_\ell}{b - a} \leqslant -\pi(\ell) \wedge \forall \ell \in L_{\mathsf{Max}}\ \frac{f_\ell(b) - x_\ell}{b - a} \geqslant -\pi(\ell)$ **then**

11 $\quad\quad\quad$ **foreach** $\ell \in L$ **do**

$\quad\quad\quad\quad$ $f_\ell := \big(\nu \in [a, b] \mapsto f_\ell(b) + (\nu - b)\frac{f_\ell(b) - x_\ell}{b - a}\big) \rhd f_\ell$

12 $\quad\quad\quad$ $b := a\,;\ stop := false$

13 $\quad\quad$ **else** $stop := true$

14 $\quad$ **until** $b = 0$ *or stop*

15 $\quad$ $r := b$

16 **return** $\boldsymbol{f}$

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

– If $\nu' > r_1$ and $\ell \in L_{\mathsf{Min}}$, we know that ℓ is non-urgent, so that $\ell \notin L'$. Therefore, by definition of σ_{Min}, $\sigma_{\mathsf{Min}}(\ell, \nu) = (r_1 - \nu + t', (\ell, \ell'))$ where $\sigma^1_{\mathsf{Min}}(\ell, \nu) = (t, (\ell, \ell^f))$ for some delay t, and $\sigma^0_{\mathsf{Min}}(\ell, r_1) = (t', (\ell, \ell'))$. 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 $\mathsf{Price}_{\mathcal{G}_{L',r_0}}(\rho) \leqslant \mathsf{Val}_{\mathcal{G}}(\ell, \nu)$.

As a consequence of this induction, we have shown that for all $\ell \in L$, and for all $\nu \in (r'', r_1)$, $\mathsf{fake}^{\sigma_{\mathsf{Min}}}_{\mathcal{G}_{L',r_0}}(\ell, \nu) \leqslant \mathsf{Val}_{\mathcal{G}}(\ell, \nu)$, which shows one inequality of (5), the other being obtained very similarly. □

Algorithm 2 implements these ideas. Each iteration of the **while** loop computes a new game in the sequence $\mathcal{G}_{L \setminus L_u, 1}$, $\mathcal{G}_{L \setminus L_u, \mathsf{left}(1)}, \dots$ described above; solves it thanks to `solveInstant`; and thus computes a new portion of $\mathsf{Val}_{\mathcal{G}}$ on an interval on the left of the current point $r \in [0, 1]$. More precisely, the vector $(\mathsf{Val}_{\mathcal{G}}(\ell, 1))_{\ell \in L}$ is first computed in line 1. Then, the algorithm enters the **while** loop, and the game $\mathcal{G}'$ obtained when reaching line 6 is $\mathcal{G}_{L \setminus L_u, 1}$. 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 $\mathsf{Val}_{\mathcal{G}'}$ that coincide with $\mathsf{Val}_{\mathcal{G}}$. It proceeds by enumerating the possible cutpoints a of $\mathsf{Val}_{\mathcal{G}'}$, starting in r, by decreasing valuations (line 8), and computes the value of $\mathsf{Val}_{\mathcal{G}'}$ in each cutpoint thanks to `solveInstant` (line 9), which yields a new piece of $\mathsf{Val}_{\mathcal{G}'}$. Then, the **if** in line 10 checks whether this new piece coincides with $\mathsf{Val}_{\mathcal{G}}$, using the condition given by Proposition 13. If it is the case, the piece of $\mathsf{Val}_{\mathcal{G}'}$ is added to f_ℓ (line 11); **repeat** is stopped otherwise. When exiting the **repeat** loop, variable b has value $\mathsf{left}(1)$. Hence, at the next iteration of the **while** loop, $\mathcal{G}' = \mathcal{G}_{L\setminus L_u,\mathsf{left}(1)}$ when reaching line 6. By continuing this reasoning inductively, one concludes that the successive iterations of the **while** loop compute the sequence $\mathcal{G}_{L\setminus L_u,1}$, $\mathcal{G}_{L\setminus L_u,\mathsf{left}(1)},\ldots$ as announced, and rebuilds $\mathsf{Val}_{\mathcal{G}}$ from them. Termination in exponential time is ensured by Proposition 20: each iteration of the **while** loop discovers at least one new cutpoint of $\mathsf{Val}_{\mathcal{G}}$, and there are at most exponentially many (note that a tighter bound on this number of cutpoints would entail a better complexity of our algorithm).

Example 1. *Figure 11 shows the value functions of the SPTG of Figure 1. Here is how the algorithm 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 price 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* $[0, 1]$.

7. Towards more complex PTGs

In [BLMR06, Rut11, HIJM13], *general* PTGs with *non-negative prices* 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 prices*. 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 $a, b \in \mathbf{N}$, consists in enhancing the locations with regions while keeping an equivalent game. This technique *can* be adapted to arbitrary weights.

Formally, let $\mathcal{G} = (L_{\mathsf{Min}}, L_{\mathsf{Max}}, L_f, L_u, \boldsymbol{\varphi}, \Delta, \pi)$ be a PTG. We define the region-PTG of $\mathcal{G}$ as $\mathcal{G}' = (L'_{\mathsf{Min}}, L'_{\mathsf{Max}}, L'_f, L'_u, \boldsymbol{\varphi}', \Delta', \pi')$ where:

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

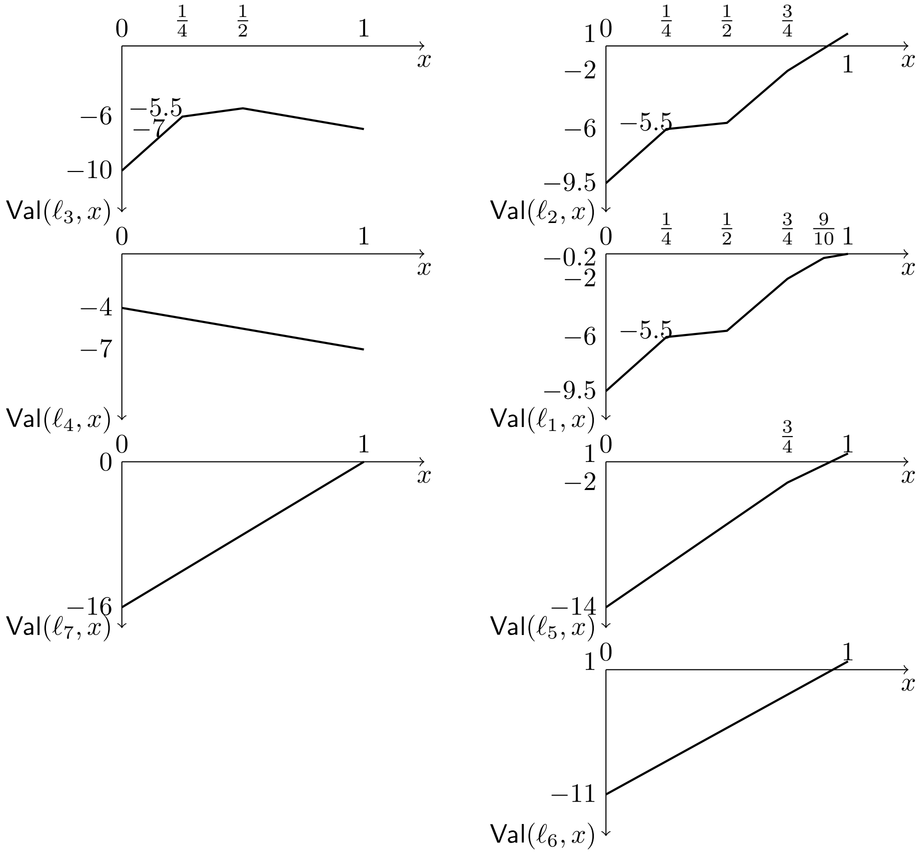

Figure 11: Value functions of the SPTG of Figure 1

- $$\Delta' = \left\{ ((\ell, I), \overline{I_g \cap I}, R, (\ell', I')) \mid (\ell, I_g, R, \ell') \in \Delta, I' = \begin{cases} I & \text{if } R = \bot \\ \{0\} & \text{otherwise} \end{cases} \right\}$$
$$\cup \left\{ ((\ell, (M_k, M_{k+1})), \{M_{k+1}\}, \bot, (\ell, \{M_{k+1}\})) \mid \ell \in L, (M_k, M_{k+1}) \in \mathsf{Reg}_{\mathcal{G}} \right\}$$
$$\cup \left\{ ((\ell, \{M_k\}), \{M_k\}, \bot, (\ell, (M_k, M_{k+1}))) \mid \ell \in L, (M_k, M_{k+1}) \in \mathsf{Reg}_{\mathcal{G}} \right\};$$

  Transitions belonging to the last two sets are called waiting transition denoted by $WaitTr$.

- $\forall (\ell, I) \in L', \pi'(\ell, I) = \pi(\ell)$; and $\forall ((\ell, I), I_g, R, (\ell', I')) \in \Delta'$, if $(\ell, I_g, R, \ell') \in \Delta$, then $\pi((\ell, I), I_g, R, (\ell', I')) = \pi(\ell, I_g, R, \ell)$, else $\pi((\ell, I), I_g, R, (\ell', I')) = 0$.

It is easy to verify that, in all configurations $((\ell, \{M_k\}), \nu)$ reachable from the null valuation, the valuation ν is M_k. More interestingly, in all configurations $((\ell, (M_k, M_{k+1})), \nu)$ reachable from the null valuation, the valuation ν is in $[M_k, M_{k+1}]$: indeed if $\nu = M_k$ (respectively, M_{k+1}), it intuitively simulates a configuration of the original game with a valuation arbitrarily close to M_k, but greater than M_k (respectively, smaller than M_{k+1}). The game can thus take transitions with guard $x > M_k$, but cannot take transitions with guard $x = M_k$ anymore.

Lemma 21. *Let $\mathcal{G}$ be a PTG, and $\mathcal{G}'$ be its region-PTG defined as before. For $(\ell, I) \in L \times \mathsf{Reg}_{\mathcal{G}}$ and $\nu \in I$, $\mathsf{Val}_{\mathcal{G}}(\ell, \nu) = \mathsf{Val}_{\mathcal{G}'}((\ell, I), \nu)$. 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. 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.

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 does not end with a waiting transition to the plays of $\mathcal{G}$. g is defined by induction on the length n of the plays so that for a play ρ of length n we have (1) $|\mathsf{Price}(\rho) - \mathsf{Price}(g(\rho))| \leqslant 2\Pi^{\mathrm{loc}}(1 - \frac{1}{2^n})\varepsilon$ and (2) there exists $\ell \in L$ and $I \in \mathsf{Reg}_{\mathcal{G}}$ such that $g(\rho)$ and ρ ends in the states ℓ and (ℓ, I) and their valuations 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 an interval 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 $\delta_{n+1} \notin WaitTr$. Let $last = \max(\{k \leqslant n \mid tr_k \notin WaitTr\})$ (with $\max \emptyset = 0$). Then, by induction, there exists $\rho' = (q_1, \nu_1) \to \dots \to (q_{last+1}, v'_{last+1})$ such that $g(\rho_{|last}) = \rho'$ (where $\rho_{|last}$ is the prefix of length $last$ of ρ), $|\mathsf{Price}(\rho_{|last}) - \mathsf{Price}(g(\rho_{|last}))| \leqslant 2\Pi^{\mathrm{loc}}(1 - \frac{1}{2^{last}})\varepsilon$ and $|\nu'_{last+1} - \nu_{last+1}| \leqslant \frac{1}{2^{last+1}}\varepsilon$. Then we choose $g(\rho) = \rho' \xrightarrow{t, \delta_{n+1}, c} (q_{n+1}, \nu'_{n+1})$ where

- if δ_{n+1} is enabled in $\mathcal{G}$ in $(q_{last+1} = q_n, \nu_n + t_{n+1})$, $t = \nu_n + t_{n+1} - \nu'_{last}$;
- otherwise, as the permissive interval of $\mathcal{G}'$ are the closure of the permissive interval 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.

Thus, in both cases, $|\nu_{n+1} - \nu'_{n+1}| \leqslant \frac{\varepsilon}{2^{n+2}}$ 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,

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

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). ρ and $g(\rho)$ ends in the locations (q, I) and q respectively.

- If ρ ends in a 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 price c. If δ belongs to $WaitTr$, then the new configuration also belongs to Max where he will make another choice. Let ρ' be the extension of ρ until the first transition δ' such that $\delta' \notin WaitTr$. $g(\rho')$ 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 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: $\sigma'_{\mathsf{Min}}(\rho)$ 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 $\forall i \leqslant k, \delta_i \in WaitTr$ 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 accepting plays of $\mathcal{G}'$ end with a transition of $WaitTr$, every accepting play ρ conforming with σ'_{Min} verifies that $g(\rho)$ conforms with σ_{Min}. Thus for every configuration s, $\mathsf{Price}_{\mathcal{G}'}(s, \sigma'_{\mathsf{Min}}) \leqslant \mathsf{Price}_{\mathcal{G}}(s, \sigma_{\mathsf{Min}}) + 2\Pi^{\mathrm{loc}}\varepsilon$. Therefore $\mathsf{Val}_{\mathcal{G}'}(s) \leqslant \mathsf{Val}_{\mathcal{G}}(s)$.

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). ρ' and ρ ends in the configuration $((q, I), \nu')$ and (q, ν).

- 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 price c and let $\tilde{\rho}$ be the extension of ρ by this choice. There exists $(t_1, \delta_1, c_1), \ldots, (t_{k+1}, \delta_{k+1}, c_{k+1})$ such that $\forall i \leqslant k, \delta_i \in WaitTr$, $\delta_{k+1} = \delta$ and $\sum_{i=1}^{k+1} t_i = \nu + t - \nu'$. Let $\rho_c = \rho' \xrightarrow{t_1, \delta_1, c_1} \ldots \xrightarrow{t_k, \delta_k, c_k} ((q, I''), \nu_k) \xrightarrow{t_{k+1}, \delta, c_{k+1}} ((q', I'), \nu_{k+1})$, 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 $\rho_c = \rho \xrightarrow{t_1, \delta_1, c_1} \ldots \xrightarrow{t_k, \delta_k, c_k} ((q, I''), \nu_k) \xrightarrow{t_{k+1}, \delta, c_{k+1}} ((q', I'), \nu_{k+1})$ such that ρ_c conforms with σ'_{Min}. We choose $\sigma_{\mathsf{Min}}(\rho) = (t, \delta)$ such that for the adequate price c, $g(\rho_c) = \rho \xrightarrow{t, \delta, c} (q', v'')$. This is possible as $t + \nu' \in I''$.

Every accepting play ρ conforming with σ_{Min} verifies $\exists \rho' \in g^{-1}(\rho)$ conforming with σ_{Min}. Thus for every configuration s, $\mathsf{Price}_{\mathcal{G}}(s, \sigma_{\mathsf{Min}}) \leqslant \mathsf{Price}_{\mathcal{G}'}(s, \sigma_{\mathsf{Min}}) + 2\Pi^{\mathrm{loc}}\varepsilon$. Therefore $\mathsf{Val}_{\mathcal{G}'}(s) \geqslant \mathsf{Val}_{\mathcal{G}}(s)$. Hence $\mathsf{Val}_{\mathcal{G}'}(s) = \mathsf{Val}_{\mathcal{G}}(s)$. □

The technique to handle resets, 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 cycle in the configuration graph visits 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 *exponential time*, and show that their value functions are cost functions with at most exponentially many cutpoints, using our techniques. In [BLMR06] the authors showed that with one-clock PTG and positive prices only we could bound the number of reset by n, the number of states, without changing the value functions. Unfortunately, the arguments to bound the number of resets do not hold for arbitrary prices, as shown by the PTG in Figure 12. We claim that $\mathsf{Val}(\ell_0) = 0$; that Min has no optimal strategy, but a family of ε-optimal strategies $\sigma^{\varepsilon}_{\mathsf{Min}}$ each with value ε; and that each $\sigma^{\varepsilon}_{\mathsf{Min}}$ requires *memory whose size depends on* ε and might *yield a play visiting at least* $1/\varepsilon$ *times the reset* between ℓ_0 and ℓ_1 (hence the number of resets cannot be bounded). For all $\varepsilon > 0$, $\sigma^{\varepsilon}_{\mathsf{Min}}$ 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 $\sigma^{\varepsilon}_{\mathsf{Min}}$, Max has two possible choices:

(i) either wait 0 time unit in ℓ_1, wait ε time units in ℓ_2, then reach ℓ_f; or

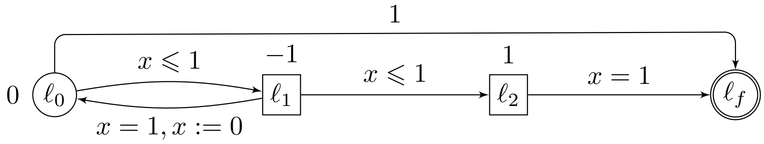

Figure 12: A PTG where the number of resets in optimal plays cannot be bounded a priori.

(ii) wait ε time unit in ℓ_1 then force the cycle by going back to ℓ_0 and wait for Min's next move.

Thus, all plays according to $\sigma^{\varepsilon}_{\mathsf{Min}}$ 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 $0 \leqslant k < \lceil 1/\varepsilon \rceil$; or of the form $\ell_0(\ell_1\ell_0)^{\lceil \frac{1}{\varepsilon} \rceil}\ell_f$. In the former case, the price of the play will be $-k\varepsilon + 0 + \varepsilon = -(k-1)\varepsilon \leqslant \varepsilon$; in the latter, $-\varepsilon(\lceil 1/\varepsilon \rceil) + 1 \leqslant 0$. This shows that $\mathsf{Val}(\ell_0) = 0$, but there is no optimal strategy as none 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 they introduce the notion of robust games. Such games requires 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 price or a price smaller than $-\kappa$. Here we require a less powerful assumption as we put this restriction only on cycles containing a reset.

Definition 5. *Given $\kappa > 0$, a κ-negative-reset-acyclic PTG (NRAPTG) is a PTG where for every state $\ell \in L$ and every cyclic finite play ρ starting and ending in $(\ell, 0)$, either $\mathsf{Price}(\rho) \geqslant 0$ or $\mathsf{Price}(\rho) < -\kappa$.*

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 price 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 price -1. The value of this PTG is 0 but no strategies for Max can achieve it because of the guard $x > 0$. As this guard is not strict anymore in the region PTG, both player 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 $k = |\mathsf{Reg}_{\mathcal{G}}|$ be the number of regions.

Lemma 22. *For all κ-NRAPTGs $\mathcal{G}$, for all $(\ell, \nu) \in \mathsf{Conf}_{\mathcal{G}}$: either $\mathsf{Val}_{\mathcal{G}}(\ell, \nu) \in \{-\infty, +\infty\}$, or*

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

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

Let $\sigma^{\neg c}_{\mathsf{Min}}$ be any memoryless strategy of Min in the reachability timed game induced by $\mathcal{G}$ such that no play consistent with $\sigma^{\neg c}_{\mathsf{Min}}$ goes twice in the same couple (location, region). If such a strategy does not exist, as the clock constraints are

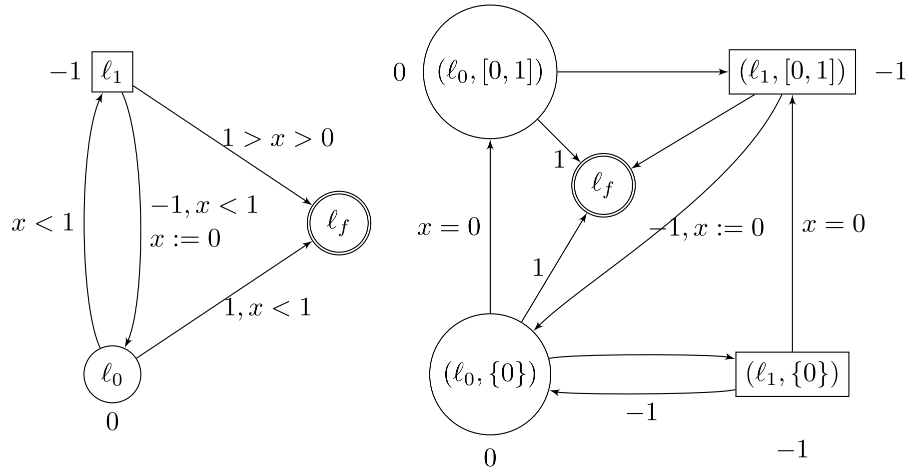

Figure 13: A −1-NRAPTG and its region PTG (some guards were removed for a better readability)

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 him and the value of $\mathcal{G}$ is $+\infty$. Let us note $\rho = \mathsf{Play}((\ell,\nu), \sigma^{\neg c}_{\mathsf{Min}}, \sigma_{\mathsf{Max}})$. By ε-optimality of σ_{Max}, $\mathsf{Price}(\rho) \geqslant \mathsf{Val}_{\mathcal{G}}(\ell,\nu) - \varepsilon$. Let $\mathsf{Price}^{\mathrm{tr}}(\rho)$ be the price of ρ due to the prices of the transitions, and $\mathsf{Price}^{\mathrm{loc}}(\rho)$ be the price due to the time elapsed in the locations of the game: $\mathsf{Price}(\rho) = \mathsf{Price}^{\mathrm{tr}}(\rho) + \mathsf{Price}^{\mathrm{loc}}(\rho)$. As there are no cycles in the game according to couples (location, region), there are at most $|L|k$ transitions, thus $\mathsf{Price}^{\mathrm{tr}}(\rho) \leqslant |L|k\Pi^{\mathrm{tr}}$. 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 $\mathsf{Price}^{\mathrm{loc}}(\rho) \leqslant |L|M\Pi^{\mathrm{loc}}$. This implies that

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

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 play $\rho = \mathsf{Play}((\ell,\nu), \sigma_{\mathsf{Min}}, \sigma_{\mathsf{Max}})$. We have that $\mathsf{Price}(\rho) \leqslant \mathsf{Val}_{\mathcal{G}}(\ell,\nu) + \varepsilon$.

We want to lower bound the prince 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 weight 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 him inside to get value $+\infty$): but then, Min could decide to cycle as long as he wants, then guaranteeing a value as low as possible, which contradicts the fact that $\mathsf{Val}(\ell,\nu) \notin \{-\infty, +\infty\}$. 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 subplay 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 $\ell' \in L_{\mathsf{Max}}$ without reset transitions, this cycle has the form $(\ell', \nu') \xrightarrow{c'} (\ell'', \nu' + t) \cdots \xrightarrow{c''} (\ell', \nu'')$ with $\nu'' \geqslant \nu'$, 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 $\mathsf{Val}_{\mathcal{G}}(\ell, \nu) + \varepsilon$, since it is still playing against an ε-optimal strategy of Min. Therefore, we can consider the subplay ρ_f of ρ where all such cycles are removes: we still have $\mathsf{Price}(\rho_f) \leqslant \mathsf{Val}_{\mathcal{G}}(\ell, \nu) + \varepsilon$.

Suppose now that ρ_f contains a cycle around a location $\ell' \in L_{\mathsf{Min}}$ 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 price of this cycle is non-negative, otherwise Min could enforce this cycle he 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 is at most $|L| - 1$ resets, $|L|$ locations of Max and $|L|k$ couples (location,region) for Min, ρ_t contains at most $|L|^2$ states 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 $\mathsf{Price}^{\mathrm{loc}}(\rho_t) \geqslant -|L|^2(|L|k + k + 1)\Pi^{\mathrm{loc}}$. Moreover, as at most $|L| - 1$ resets are taken in ρ_f and that the game is bounded by M, $\mathsf{Price}^{\mathrm{loc}}(\rho_f) \geqslant -|L|M\Pi^{\mathrm{loc}}$. This implies that

$$\mathsf{Val}_{\mathcal{G}}(\ell, \nu) + \varepsilon \geqslant \mathsf{Price}^{\mathrm{loc}}(\rho_f) + \mathsf{Price}^{\mathrm{tr}}(\rho_t) \geqslant -|L|M\Pi^{\mathrm{loc}} - |L|^2(|L|k + k + 1)\Pi^{\mathrm{tr}} .$$

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 then if the generated cycle has positive price, 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 stop Min to reach an accepting state. On the contrary if the cycle has negative price, then by definition of a κ-NRAPTG, this price 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 23. *For all $\kappa > 0$, the value of a κ-NRAPTG can be computed by solving $k = \lceil \kappa \times 2n(\nu^{sup} - \nu_{inf}) \rceil$ PTGs without resets and using the same set of guards where ν^{sup} and ν_{inf} are the upper and lower bound of the value of the game given by Lemma 22. Moreover, from ε-optimal strategies on those k games, we can build $k\varepsilon$-optimal strategies in the original game.*

With this, we can conclude:

Theorem 24. *Let $\kappa > 0$ and $\mathcal{G}$ be a κ-NRAPTG. Then for every location $q \in Q$, the function $v \mapsto Val_{\mathcal{G}}(q, \nu)$ is computable in EXPTIME and are piecewise-affine functions with at most an exponential number of cutpoints. Moreover, for* *every $\varepsilon > 0$, there exists (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 25. *The value functions of all one-clock PTGs are cost functions with at most exponentially many cutpoints.*

Proof. Let $\mathcal{G}$ be a one-clock PTG. Let us replace all transitions $(\ell, g, \top, \ell')$ resetting the clock by $(\ell, g, \bot, \ell'')$, where ℓ'' is a new final location with $\varphi_{\ell''} = \mathsf{Val}_{\mathcal{G}}(\ell, 0)$—observe that $\mathsf{Val}_{\mathcal{G}}(\ell, 0)$ exists even if we cannot compute it, so this transformation is well-defined. This yields a reset-acyclic PTG $\mathcal{G}'$ such that $\mathsf{Val}_{\mathcal{G}'} = \mathsf{Val}_{\mathcal{G}}$. $\square$

8. 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 exponential time for the special case of simple games. This complexity result is comparable with previously obtained results in the case of non-negative weights only [HIJM13, Rut11], even though we follow different paths to prove termination and correction (due to the presence of negative prices). 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.

[AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.

[ALTP04] Rajeev Alur, Salvatore La Torre, and George J. Pappas. Optimal paths in weighted timed automata. *Theoretical Computer Science*, 318(3):297–322, 2004.

[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.

[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.

[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.

[BBM06] Patricia Bouyer, Thomas Brihaye, and Nicolas Markey. Improved undecidability results on weighted timed automata. *Information Processing Letters*, 98(5):188–194, 2006.

[BBR05] Thomas Brihaye, Véronique Bruyère, and Jean-François Raskin. On optimal timed strategies. In *Proceedings of the Third international conference on Formal Modeling and Analysis of Timed Systems (FORMATS'05)*, volume 3829 of *Lecture Notes in Computer Science*, pages 49–64. Springer, 2005.

[BCFL04] Patricia Bouyer, Franck Cassez, Emmanuel Fleury, and Kim G. Larsen. Optimal strategies in priced timed game automata. In *Proceedings of the 24th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'04)*, volume 3328 of *Lecture Notes in Computer Science*, pages 148–160. Springer, 2004.

[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 *LNCS*, pages 128–137. Springer, 2009.

[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, 2014*, pages 283–292, 2014.

[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.

[BGH$^+$15] Thomas Brihaye, Gilles Geeraerts, Axel Haddad, Engel Lefaucheux, 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.

[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.

[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.

[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.

[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.

[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.

[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 Concurrecy Theory (CONCUR'01)*, volume 2154 of *Lecture Notes in Computer Science*, pages 536–550. Springer, 2001.

[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.

[Mar75] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.

[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.

[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.