Almost Optimal Strategies in One Clock Priced Timed Games

Patricia Bouyer[1], Kim G. Larsen[2],
Nicolas Markey[1], and Jacob Illum Rasmussen[2]

[1] Lab. Spécification & Vérification, CNRS & ENS Cachan, France
{bouyer,markey}@lsv.ens-cachan.fr
[2] Department of Computer Science, Aalborg University, Denmark
{kgl,illum}@cs.aau.dk

Abstract. We consider timed games extended with cost information, and prove computability of the optimal cost and of ε-optimal memoryless strategies in timed games with one clock. In contrast, this problem has recently been proved undecidable for timed games with three clocks.

1 Introduction

An interesting direction of real-time model checking that has recently received substantial attention is to extend and re-target timed automata technology towards optimal scheduling and planning [1, 15, 9]. In particular, as part of this effort, the notion of priced timed automata [6, 5] has been promoted as a useful extension of the classical model of timed automata [4]. In this extended model each location q is associated with a cost c_q giving the cost of a unit of time spent in q. Thus, each run of a priced timed automaton has an accumulated cost, based on which a variety of optimization problems may be formulated.

Several of the established results concerning priced timed automata are concerned with reachability questions. In [3] cost-bounded reachability was shown decidable. [6] and [5] independently show computability of the cost-optimal reachability for priced (or weighted) timed automata using different adaptations of the so-called region technique. In [13, 15] the notion of priced zone is developed allowing efficient implementation of cost-optimal reachability as witnessed by the competitive tool UPPAAL Cora [16]. Also the problem of computing optimal *infinite* schedules (in terms of minimal limit-ratios) has been shown computable [8]. Finally cost-optimal reachability has been shown decidable in a setting with multiple cost-variables [14].

In this paper we consider the more challenging problem of the computation of *cost-optimal winning strategies* for priced timed *game* automata, *i.e.* a game where the controller tries to win at minimal cost and opponent tries to maximize the cost. Consider the priced timed game with the single clock x depicted in Fig. 1. Here the (circle) locations c_1 and c_2 are controllable whereas (square) locations u_1 and u_2 are uncontrollable with cost-rates being 3, 4, 1 and 1, respectively. All four locations have $x \leq 1$ as invariant. Besides transitions between these four locations, additional transitions are indicated to (triangle) locations for which the optimal costs of winning (for any value of x) are assumed to have already been computed (we call those cost functions *outside cost functions* in the sequel). Obviously, c_1 and c_2 have winning strategies for all values of x by uniformly exiting to their respective outside locations (triangle), c_1^{out} and c_2^{out}. However, this strategy is, clearly, suboptimal for both locations. Alternatively, consider the superior strategy for c_2 depicted in Fig. 2. that guarantees cost no larger than depicted in the corresponding cost function. Then it can be shown that this strategy guarantees the optimal cost.

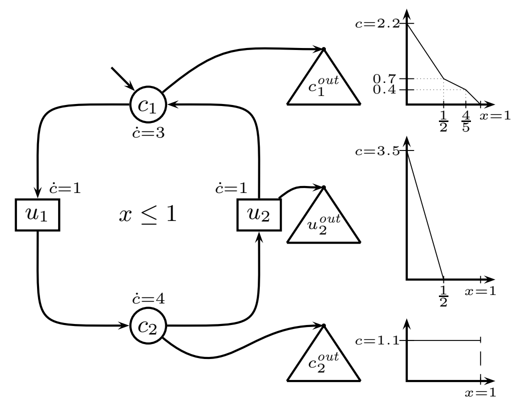

Fig. 1. Sample PTGA with outside cost functions.

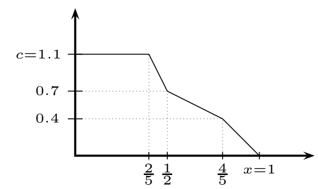

$$\sigma(c_2, x) = \begin{cases} c_2^{out} & \text{if } 0 \leq x < 2/5 \\ c_2 & \text{if } 2/5 \leq x < 1/2 \\ u_2 & \text{if } 1/2 \leq x \leq 1 \end{cases}$$

Fig. 2. An optimal strategy in c_2, and the associated cost function.

In [12] the problem of computing cost-optimal winning strategies has been studied and shown computable for *acyclic* priced timed games. Furthermore, in [11], it is proven that computing optimal winning strategies for one-clock PTGA with stopwatch cost (*i.e.* cost are either zero or one) is decidable. [2] and [10] provide partial solutions to the general case of non-acyclic games: under the assumption of certain non-Zenoness behaviour of the underlying priced timed automata it is shown that it suffices only to consider strategies guaranteed to win within some given number k of steps, or alternatively to unfold the given game k times and reduce the problem to solving an acyclic game. To see how restricted these results are, it may be observed that the priced timed game in Fig. 1 does not belong to any of the above classes. In fact, in [11] it has recently been shown that the problem of determining cost-optimal winning strategies for priced timed games is not computable. Most recently, it has been shown that this negative result holds even for priced timed (game) automata with no more than three clocks [7].

In this paper we completely solve the computation of cost-optimal winning strategies for arbitrary priced timed (game) automata *with one clock*: we offer an algorithm for computing optimal costs, explain why optimal strategies need not always exist, whereas memoryless ε-optimal strategies exist and can be computed.

2 Definitions

We write x for the (unique) clock variable, and $\mathcal{X} = \{x\}$. A clock constraint for clock x is an expression of the form $x \in I$ where I is an interval over the reals with integer (or infinite) bounds which can have strict or non-strict bounds. As a shortcut, we may use expressions like $x \geq 5$ instead of $x \in [5, +\infty[$. The set of all clock constraints is denoted $\mathcal{B}(\mathcal{X})$. That a valuation $v\colon \mathcal{X} \to \mathbb{R}_+$ satisfies a clock constraint g is defined in a natural way (v satisfies $x \in I$ whenever $v(x) \in I$), and we then write $v \models g$. We denote by v_0 the valuation that assigns zero to clock x, by $v + t$ ($t \in \mathbb{R}_+$) the valuation that assigns $v(x) + t$ to $x \in \mathcal{X}$.

A *cost function* is a piecewise affine function $f\colon \mathbb{R}_+ \to \mathbb{R}_+ \cup \{+\infty\}$ with negative slopes. We also require that if $\{+\infty\} \in f((n, n+1))$ for some integer n, then $f((n, n+1)) = \{+\infty\}$, and that f is continuous over all intervals $(n, n+1)$. We write CF for the set of all cost functions.

We define an extended notion of priced timed games, with outside cost functions and urgent locations. Those extra features will be needed throughout the proof. A 1-*clock priced timed game with outside cost functions* (PTG_f for short) is a tuple $G = (Q_c, Q_u, Q_f, Q_{\mathrm{urg}}, Q_{\mathrm{init}}, f_{\mathrm{goal}}, T, \eta, P)$ where

- Q_c is a finite set of *controllable locations*, Q_u is a finite set of *uncontrollable locations*. Those sets are disjoint, and we define $Q = Q_c \cup Q_u$;
- Q_f is the set of final locations (it is disjoint from Q).
- $Q_{\mathrm{urg}} \subseteq Q_u$ indicates urgent uncontrollable locations;
- $Q_{\mathrm{init}} \subseteq Q$ is the set of *initial* locations;
- $f_{\mathrm{goal}}\colon Q_f \to \mathrm{CF}$ assigns to each final location a cost function;
- $T \subseteq Q \times \mathcal{B}(\mathcal{X}) \times 2^{\mathcal{X}} \times (Q \cup Q_f)$ is the set of *transitions*;
- $\eta\colon Q \to \mathcal{B}(\mathcal{X})$ defines the *invariants* of each location;
- $P\colon Q \cup T \to \mathbb{N}$ is the *cost (or price) function*.

Standard (1-clock) priced timed games [2, 10] are PTG_f with $Q_{\mathrm{urg}} = \varnothing$ and, for any $q \in Q_f$, $f_{\mathrm{goal}}(q)(\mathbb{R}_+) = \{0\}$ or $\{+\infty\}$.

In the following, G will always refer to a PTG_f, and we will not always rewrite the corresponding tuple. Similarly, G' will denote a PTG_f whose components are "primed".

We assume (w.l.o.g., see [6]) that the clock is bounded, *i.e.*, there exists an integer M such that for every location $q \in Q$, $\eta(q) \Rightarrow x \leq M$.

Let G be a PTG_f. The semantics of G is given as a labeled timed transition system $\mathcal{T} = (S, S_{\mathrm{init}}, \to)$ where $S \subseteq (Q \cup Q_f) \times \mathbb{R}_+$ is the set of states[3], $S_{\mathrm{init}} = Q_{\mathrm{init}} \times \{v_0\}$ is the set of initial states, and the transitions relation $\to\, \subseteq S \times \mathbb{R}_+ \times S$ is defined as:

1. *(discrete transition)* $(q, v) \xrightarrow{c} (q', v')$ if $q \notin Q_f$ and there exists $(q, g, R, q') \in T$ such that $v(x) \models g$, $v' = [R \leftarrow 0]v$, $v'(x) \models \eta(q')$, and $c = P(q, g, R, q')$;
2. *(delay transition)* $(q, v) \xrightarrow{c} (q, v + t)$ if $q \notin Q_{\mathrm{urg}} \cup Q_f$, and $\forall 0 \leq t' \leq t$, $v + t' \models \eta(q)$, and $c = t \cdot P(q)$.

[3] Formally, $S \subseteq (Q \cup Q_f) \times (\mathbb{R}_+)^{\mathcal{X}}$, but we identify v with $v(x)$ here.

A *run* of G is a (finite) path in the underlying transition system. Given $T, U \subseteq S$, we write[4] $\mathsf{Run}_G(T, U)$ for the set of runs of G issued from $t \in T$ and ending in $u \in U$. Given a run ϱ and a position $v \in \varrho$ along that run, the prefix of ϱ ending in v is denoted by $\varrho_{|v}$. A run is *maximal* if either it is infinite, or no discrete transition is possible (even after a delay transition). A maximal run is *accepting* if it is finite and ends in a final location. Let $\varrho = s_0 \xrightarrow{c_0} s_1 \xrightarrow{c_1} \cdots \xrightarrow{c_{n-1}} s_n$ be a run. Its cost, denoted $\mathsf{cost}(\varrho)$, is either $\sum_{i=0}^{n-1} c_i$ if ϱ is not accepting, or $\sum_{i=0}^{n-1} c_i + f_{\text{goal}}(q_n)(v_n(x))$, where $(q_n, v_n) = s_n$ if ϱ is accepting. An accepting run is *winning* if it has finite cost.

Example. Reconsider the example depicted in Fig. 1. Here, a sample winning run is $\varrho = (c_1, 0) \xrightarrow{0} (u_1, 0) \xrightarrow{0.4} (u_1, 0.4) \xrightarrow{0} (c_2, 0.4) \xrightarrow{0.4} (c_2, 0.5) \xrightarrow{0} (c_2^{out}, 0.5)$ which has cost $\mathsf{cost}(\varrho) = 0.4 \times 1 + 0.1 \times 4 + f_{\text{goal}}(c_2^{out})(0.5) = 1.9$. $\square$

A *strategy* is then a function $\sigma\colon \mathsf{Run}_G(Q \times \mathbb{R}_+, Q_c \times \mathbb{R}_+) \to \{\lambda\} \cup Q \cup Q_f$. Informally, a strategy tells in all controllable locations, what has to be done, and the special symbol λ indicates to delay. A strategy σ is *memoryless* if $\sigma(\varrho) = \sigma(\varrho')$ as soon as ϱ and ϱ' end in the same state.

Let σ be a strategy in G, and ϱ_0 a run in G ending in (q_0, x_0). A run $\varrho = (q_0, x_0) \xrightarrow{c_0} (q_1, x_1) \xrightarrow{c_1} \cdots \xrightarrow{c_{n-1}} (q_n, x_n)$ is a (σ, ϱ_0)-run if for all delay- (or discrete-) transitions $(q_i, x_i) \xrightarrow{c_i} (q_{i+1}, x_{i+1})$ where $q_i \in Q_c$, we have

- $\forall x \in [x_i, x_{i+1}[,\ \sigma(\varrho_0 \cdot \varrho_{|x}) = \lambda$,
- $\sigma(\varrho_0 \cdot \varrho_{|x_i}) = q_{i+1}$.

where $\varrho_0 \cdot \varrho$ denotes the (usual) concatenation. In that case, we say that ϱ is *compatible* with σ after ϱ_0 (or that it is an *outcome* of σ after ϱ_0). We write $\mathsf{Run}_{G,\sigma}(\varrho_0, U)$ for the set of such runs ending in U.

A strategy σ is said *accepting* after (run) ϱ_0 whenever all maximal runs in $\mathsf{Run}_{G,\sigma}(\varrho_0)$ are accepting. If a strategy is not accepting from ϱ_0, we set its cost in G after ϱ_0, $\mathsf{Cost}_G(\sigma, \varrho_0)$, to $+\infty$. Otherwise its cost in G after ϱ_0 is given as: $\mathsf{Cost}_G(\sigma, \varrho_0) = \sup\{\mathsf{cost}(\varrho) \mid \varrho \in \mathsf{Run}_{G,\sigma}(\varrho_0, Q_f \times \mathbb{R}_+)\}$. Obviously, for any two runs ϱ_0 and ϱ_1 ending in (q, x), the sets $\{\mathsf{Cost}_G(\sigma, \varrho_0) \mid \sigma \text{ strategy in } G\}$ and $\{\mathsf{Cost}_G(\sigma, \varrho_1) \mid \sigma \text{ strategy in } G\}$ are equal. An accepting strategy σ after ϱ_0 is *winning* if $\mathsf{Cost}_G(\sigma, \varrho_0)$ is finite. We define for every state s of G, the *optimal cost* of winning from s as $\inf\{\mathsf{Cost}_G(\sigma, \varrho_0) \mid \sigma \text{ strategy in } G\}$ for some run ϱ_0 ending in s. We denote it $\mathsf{OptCost}_G(s)$. If $\mathsf{OptCost}_G(s) < +\infty$, the state s is said *winning* in G. In that case, for every $\varepsilon > 0$, for every run ϱ_0 ending in s, there exists a winning strategy σ s.t. $\mathsf{OptCost}_G(s) \le \mathsf{Cost}_G(\sigma, \varrho_0) < \mathsf{OptCost}_G(s) + \varepsilon$, and we say that σ is *ε-optimal* from ϱ_0. A strategy σ is *optimal* from ϱ_0 if $\mathsf{Cost}_G(\sigma, \varrho_0) = \mathsf{OptCost}_G(s)$ where ϱ_0 ends in state s.

A strategy σ in G is *(ε, N)-acceptable* (with $\varepsilon > 0$, and $N \in \mathbb{N}$) whenever: (1) it is memoryless, (2) it is ε-optimal, (3) there exist N (consecutive) intervals $(I_i)_{1 \le i \le N}$ partitioning $[0, 1]$ such that for every location q, for every $1 \le i \le N$, for every integer $\alpha < M$, the function $x \mapsto \mathsf{Cost}_G(\sigma, (q, x))$ is affine on every interval $\alpha + I_i$, and the function $x \mapsto \sigma(q, x)$ is constant on $\alpha + I_i$.

[4] In the sequel, we might omit the subscripts G when they are clear from the context.

3 Main result

The main result of this paper is that optimal cost is computable and that almost-optimal memoryless strategies always exist and can be effectively computed. This is summarized by the following theorem:

Theorem 1. *Let G be a PTG_f. Then for every location q in G, the function $x \mapsto \mathsf{OptCost}_G((q,x))$ is computable and piecewise-affine. Moreover, for every $\varepsilon > 0$, there exists (and we can effectively compute) a strategy σ in G such that σ is memoryless and ε-optimal in every state.*

We will even prove a stronger result, which is that there exists $N \in \mathbb{N}$ such that for every $\varepsilon > 0$, we can effectively compute an (ε, N)-acceptable strategy σ. The rest of this paper is devoted to a proof of this result.

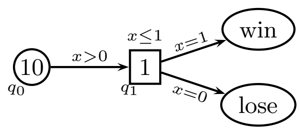

Fig. 3. A game with no optimal strategy

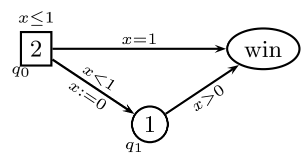

Fig. 4. A game where optimal strategies require memory

There are PTG_f for which no optimal strategies exist, as exemplified by Fig. 3: from q_0, the optimal cost is 1, but a winning strategy consists in delaying in q_0 for some duration $\delta > 0$, yielding a cost of $1 + 9\delta$. This is why we compute, in the general case, ε-optimal strategies. In the same way, as witnessed by Fig 4, it might be the case that optimal strategies exist but require some amount of memory: in the example of Fig 4, state $(q_0, x = 0)$ is winning with optimal cost 2, but no memoryless strategy can achieve that cost for sure.

4 Simplifying Transformations

In this section, we first explain how to restrict to simpler games while preserving the same optimal costs, and we then show how we can inductively compute optimal cost on those simpler games. We also explain how to compute almost-optimal strategies for those simpler games, and how to "lift" those strategies to the original game.

Our transformations consist in two steps: (i) we restrict to PTG_f where the clock is bounded by 1 (denoted $[0,1]$-PTG_f) (Section 4); (ii) we restrict to a $[0,1]$-PTG_f without resetting transition (Section 4). For each transformation, we prove that:

– the optimal cost in each state of the original game is identical to the optimal cost in some corresponding state in the transformed game,

– we can derive an ε-optimal strategy in the original game from some ε'-optimal strategy in the transformed game.

Section 5 is then devoted to computing the optimal cost and an almost-optimal strategy in the simpler game. For the sake of simplicity, we assume here that there are no discrete costs on transitions. A slight adaptation of the transformation for removing resets can be given for handling discrete costs as well.

Restricting to a PTG$_f$ bounded by 1. The idea of this construction is to reset the clock each time it reaches 1, and to record in the discrete structure what should be the real integer part of the value of the clock (the clock will only store the fractional part of its real value).

Let G be a PTG$_f$. We build another PTG$_f$ G' such that for every $q' \in Q'$, $\eta'(q')$ implies $0 \leq x \leq 1$, and G' is correct for computing optimal cost, in a sense which will be made clear later.

As we have assumed that PTG$_f$ are bounded, we set M the constant bounding G, and we define:

$$\begin{cases} Q'_x &= \{q_{[\alpha,\alpha+1]} \mid q \in Q_x \text{ and } 0 \leq \alpha < M\} \text{ for every } x \in \{c, u, f, \mathrm{urg}\} \\ Q'_{\mathrm{init}} &= \{q_{[0,1]} \mid q \in Q_{\mathrm{init}}\} \end{cases}$$

The set of transitions T' is composed of the following transitions (if g is a guard, $g - \alpha$ denotes the same guard translated by $-\alpha$):

$$\begin{cases} q_{[\alpha,\alpha+1]} \xrightarrow{(g-\alpha)\cap(0\leq x<1)} q'_{[\alpha,\alpha+1]} & \text{if } (q \xrightarrow{g} q') \in T \text{ and } \alpha+1 < M \\ q_{[M-1,M]} \xrightarrow{(g-\alpha)\cap(0\leq x\leq 1)} q'_{[M-1,M]} & \text{if } (q \xrightarrow{g} q') \in T \\ q_{[\alpha,\alpha+1]} \xrightarrow[x:=0]{(g-\alpha)\cap(0\leq x<1)} q'_{[0,1]} & \text{if } (q \xrightarrow[x:=0]{g} q') \in T \text{ and } \alpha+1 < M \\ q_{[M-1,M]} \xrightarrow[x:=0]{(g-\alpha)\cap(0\leq x\leq 1)} q'_{[0,1]} & \text{if } (q \xrightarrow[x:=0]{g} q') \in T \\ q_{[\alpha-1,\alpha]} \xrightarrow[x:=0]{x=1} q_{[\alpha,\alpha+1]} & \text{if } 0 < \alpha < M \end{cases}$$

The invariant η' is defined by $\eta'(q_{[\alpha,\alpha+1]}) = (0 \leq x \leq 1) \wedge (\eta(q) - \alpha)$ if $q \in Q$. The cost function P' is defined by $P'(q_{[\alpha,\alpha+1]}) = P(q)$. The function f'_{goal} is defined by $f'_{\mathrm{goal}}(q_{[\alpha,\alpha+1]})(x) = f_{\mathrm{goal}}(q)(x+\alpha)$ for every $0 \leq x \leq 1$.

Note that all guards and invariants of G' are included in $[0,1]$, we say that G' is a $[0,1]$-PTG_f .

We define f the function which maps every state (q,x) of G onto the state $(q_{[\alpha,\alpha+1]}, x-\alpha)$ of G' such that $0 \leq x - \alpha \leq 1$ and $x < M$ integer implies $x = \alpha$. We now state the following correctness result.

Proposition 2. *For every state (q,x) in G, $\mathsf{OptCost}_G(q,x) = \mathsf{OptCost}_{G'}(f(q,x))$. Moreover, for every $\varepsilon > 0$ and $N \in \mathbb{N}$, given an (ε,N)-acceptable strategy in G', we can compute an (ε,N)-acceptable strategy in G, and vice-versa.*

Removing resetting transitions from SCCs. We have restricted to games with a single clock. A strong property of this model is that each time a resetting transition is taken, then the *very same state* is visited (because the valuation is each time v_0). The construction for removing resetting transitions takes advantage of this property.

Let G be a PTG_f with n resetting transitions. From the previous reduction, we may assume that all the invariants and guards in G imply that $0 \le x \le 1$. We build a PTG_f G', made of $n+1$ copies of G, such that no strongly connected component (SCC for short) of G' contains a resetting transition.

We thus define $Q'_c = Q_c \times \{0, ..., n\}$, $Q'_u = Q_u \times \{0, ..., n\}$, and $Q'_f = (Q_f \times \{0, ..., n\}) \cup \{r\}$. A location $(q, i) \in Q'_u$ is urgent iff $q \in Q_{\mathrm{urg}}$. We let $Q'_{\mathrm{init}} = Q_{\mathrm{init}} \times \{0\}$. The outside cost functions are given by $f'_{\mathrm{goal}}((q, i)) = f_{\mathrm{goal}}(q)$, and $f_{\mathrm{goal}}(r) = +\infty$. The invariant is given by $\eta'((q, i)) = \eta(q)$ for $q \in Q$. Transitions are defined as follows:

$$\begin{cases} ((q,i) \xrightarrow{g} (q',i)) \in T' & \text{if } (q \xrightarrow{g} q') \in T \text{ and } i \le n \\ ((q,i) \xrightarrow[x:=0]{g} (q',i+1)) \in T' & \text{if } (q \xrightarrow[x:=0]{g} q') \in T \text{ and } i < n \\ ((q,n) \xrightarrow{g} r) \in T' & \text{if } (q \xrightarrow[x:=0]{g} q') \in T \text{ and } i = n \end{cases}$$

Last, we set $P'((q, i)) = P'(q)$ for every $q \in Q$, and the price of each transition of T' defined above is the price of the corresponding transition in T.

Proposition 3. *For every state (q, x) in the game G, $\mathsf{OptCost}_G((q, x))$ equals $\mathsf{OptCost}_{G'}(((q, 0), x))$. Moreover, for every $\varepsilon' > 0$ and $N' \in \mathbb{N}$, given an (ε', N')-acceptable strategy in G', we can compute a $(2\varepsilon', N')$-acceptable strategy in G.*

We have thus reduced our problem to computing optimal cost and almost-optimal winning strategies in G'. In G', this can be done by first computing it in the nth copy of G, and then in the $(n-1)$th copy of G, etc.

5 Computing almost-optimal strategies

We have restricted our problem to $[0, 1]$-PTG_f without resets. We can also easily restrict to such PTG_f containing only one SCC: if we can compute the optimal costs and an (ε, N)-acceptable strategy on an SCC, we will be able to handle the general case by working first on the deepest SCC, and then replace it by the corresponding outside function (and an (ε, N)-acceptable strategy).

Thus, we now assume that we only work on a $[0, 1]$-PTG_f without resets and based on an SCC. We prove the following result, which will imply Theorem 1.

Theorem 4. *Let G be a $[0, 1]$-PTG_f without reset such that $(Q_c \cup Q_u, T)$ is an SCC (or contains only one location). Then:*

H1. *$\mathsf{OptCost}_G(q, x)$ is computable for every $q \in Q$ and every $x \in [0, 1]$;*
H2. *for every location $q \in Q$, $x \in [0, 1] \mapsto \mathsf{OptCost}_G(q, x)$ is a cost function whose finitely many segments either have slope $-c$ where $c \in P(Q)$, or are fragments of the outside cost functions of G;*

H3. there exists an integer N such that, for any $\varepsilon > 0$, we can compute an (ε, N)-acceptable strategy in G for every $q \in Q$ and every $x \in [0,1]$.

The rest of this section is devoted to the proof of this theorem, which is by induction on the number of non-urgent locations in G. First we prove the base case of the induction, that is when the game is only composed of urgent locations, or of a single controllable location.

- Proving properties H1 and H2 in the case where G contains only one location is handled straightforwardly, by combining the outside cost functions of G with the cost rate of the location. Property H3 requires more care. Let q be a (controllable) location with a bunch of outside cost functions $\{f_{\text{goal}}(q') \mid q' \in Q_f\}$. Define the function $s\colon x \to \min\{f_{\text{goal}}(q', x) \mid q' \in Q_f\}$. Then $\mathsf{OptCost}_G(q, x) = \inf_{x \leq x' \leq 1} P(q) \cdot (x' - x) + s(x')$. Let $\varepsilon > 0$. We then define the strategy σ as follows:

$$\sigma(q,x) = \begin{cases} q' \text{ if } \mathsf{OptCost}_G(q,x) = f_{\text{goal}}(q')(x) \\ \lambda \text{ if } \mathsf{OptCost}_G(q,x) < s(x) \text{ and either } s(1) < +\infty \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{or } x \leq 1 - \varepsilon/(2P(q)) \\ q' \text{ if } \mathsf{OptCost}_G(q,x) < s(x),\ s(1) = +\infty,\ 1 - \varepsilon/(2P(q)) < x < 1, \\ \quad \text{and } \lim_{x \to 1^-} f_{\text{goal}}(q')(x) = \lim_{x \to 1^-} s(x) \end{cases}$$

  It is not difficult to check that σ is (ε, N)-acceptable for some N which is independent of ε.
- The case where G contains an SCC with only urgent (thus uncontrollable) locations is also straightforward, since the opponent can force the game to never reach a final location, and the optimal cost is then $+\infty$. If the game is composed of a single urgent location, then this is also easy.

We now assume that G is an SCC composed of at least two locations, n of which are non-urgent. We select one of the non-urgent locations having least cost, and denote it with $q_{\min}$, and, depending on the nature (controllable or not) of $q_{\min}$, we explain how we prove that Theorem 4 holds for G if it holds for SCCs having at most $(n-1)$ non-urgent locations.

Case: $q_{\min}$ is controllable. For handling this case, we will prove that the rough intuition that there is no need to delay twice in $q_{\min}$, but we better delay longer in $q_{\min}$ is indeed correct.

From the game G, we construct a game G', made of two copies of G, such that each SCC of the new game contains one location less (see Fig. 5). We define $Q'_c = (Q_c \setminus \{q_{\min}\}) \times \{0,1\} \cup \{q_{\min}\}$, $Q'_u = Q_u \times \{0,1\}$, $Q'_f = Q_f \times \{0,1\} \cup \{r\}$, $Q'_{\text{urg}} = Q_{\text{urg}} \times \{0,1\}$, $Q'_{\text{init}} = Q_{\text{init}} \times \{0\}$, $f'_{\text{goal}}((q,i)) = f_{\text{goal}}(q)$ if $q \in Q_f$, and $f'_{\text{goal}}(r) = +\infty$, $\eta'((q,i)) = \eta(q)$, $\eta'(q_{\min}) = \eta(q_{\min})$, $P'((q,i)) = P(q)$ for every $(q,i) \in Q'_c \cup Q'_u$. The set of transitions is

$$\begin{aligned} T' = {} & \{(q,i) \xrightarrow{g,R} (q',i) \mid q \xrightarrow{g,R} q', \text{and } q, q' \neq q_{\min}\} \\ & \cup \{(q,0) \xrightarrow{g,R} q_{\min},\ (q,1) \xrightarrow{g,R} r \mid (q \xrightarrow{g,R} q_{\min}) \in T\} \\ & \cup \{q_{\min} \xrightarrow{g,R} (q',1) \mid (q_{\min} \xrightarrow{g,R} q') \in T\}. \end{aligned}$$

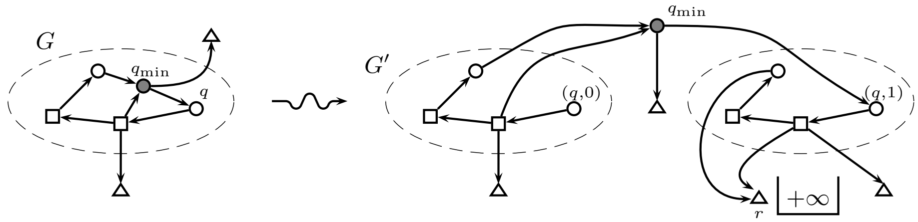

Fig. 5. Case $q_{\min}$ (in grey) controllable

We prove the following lemma, which establishes properties H1 and H2.

Lemma 5. *For every* $(q, x) \in (Q \setminus \{q_{\min}\}) \times [0, 1]$*, we have* $\mathit{OptCost}_G(q, x) = \mathit{OptCost}_{G'}((q, 0), x)$*. For every* $x \in [0, 1]$*,* $\mathit{OptCost}_G(q_{\min}, x) = \mathit{OptCost}_{G'}(q_{\min}, x)$*.*

It remains to prove property H3. We fix the integer N' for G'. We fix some $\varepsilon > 0$, and take $\varepsilon' = \frac{\varepsilon}{3}$. We take σ' an (ε', N')-acceptable strategy in G'. We then define σ as follows:

$$\sigma(q, x) = \begin{cases} \sigma'((q, 1), x) & \text{if } \mathsf{Cost}_{G'}(\sigma', ((q, 1), x)) \le \mathsf{OptCost}_{G'}(q_{\min}, x) \\ \sigma'((q, 0), x) & \text{otherwise} \end{cases} \tag{1}$$

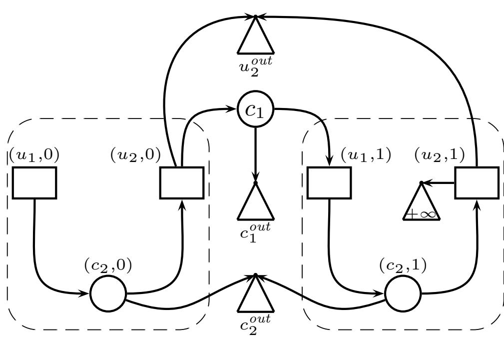

Fig. 6. Running example after unwinding.

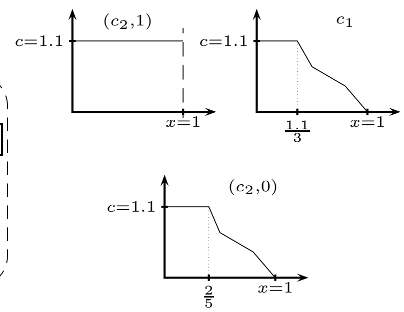

Fig. 7. Optimal costs.

Example. Returning to the running example of Fig. 1 with u_1 and u_2 urgent, performing the above transformation with respect to c_1 gives the PTG$_f$ depicted in Fig. 6. The optimal cost functions are depicted in Fig. 7 and the resulting winning strategy for c_2 is, according to (1), the strategy of $(c_2, 1)$ when $x \le \frac{1.1}{3}$ and $(c_2, 0)$ otherwise. □

Obviously, the strategy σ is memoryless. We need to establish that the function $x \mapsto \mathsf{Cost}_G(\sigma, (q, x))$ consists of at most N pieces, and that σ is ε-optimal.

Proposition 6. *Strategy* σ *is winning and there exists a fixed (independant of* ε*) integer* N *such that* σ *is* (ε, N)*-acceptable.*

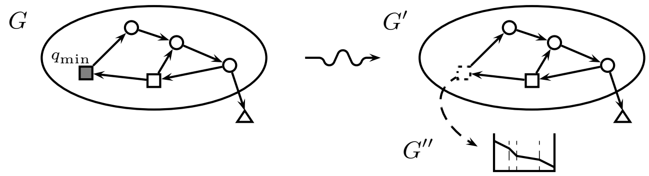

Fig. 8. When it is uncontrollable, $q_{\min}$ is made urgent (in dash line here).

Case: $q_{\min}$ is uncontrollable. The intuition is that the opponent will prefer delays in other locations than $q_{\min}$ whenever possible. We attempt to enforce this by a transformation of the game where location $q_{\min}$ is urgent, as depicted in Fig. 8. Formally, given a $[0,1]$-PTG$_f$ without resets G, we define G' with $Q'_{\text{urg}} = Q_{\text{urg}} \cup \{q_{\min}\}$ and $Q'_u = Q_u \backslash \{q_{\min}\}$.

Obviously enough, since we restrict the possible moves for the opponent in G', we have for every state (q,x), $\mathsf{OptCost}_{G'}(q,x) \leq \mathsf{OptCost}_G(q,x)$.

However, the converse inequality is not correct over $[0,1]$, and we will need a more complex construction to handle this case. We now explain how to iteratively compute the optimal costs in G. Fig. 9 gives an overview of the computation described below.

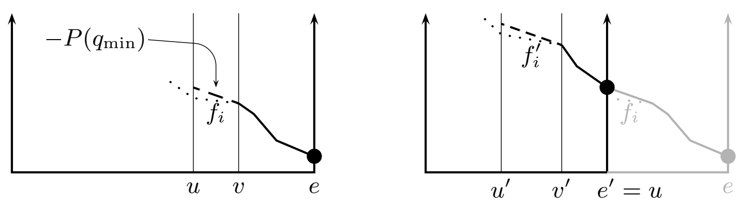

Fig. 9. Successive computations when $q_{\min}$ is uncontrollable

Clearly, we can compute $\mathsf{OptCost}_G(q_{\min},1)$ (indeed, $\mathsf{OptCost}_G(q_{\min},1) = \mathsf{OptCost}_{G'}(q_{\min},1)$, since when $x=1$, time cannot elapse any more and the same moves are available in G' and in G). This initializes our iterative computation.

Now, assume we can compute $\mathsf{OptCost}_G(q_{\min},e)$ for some $e \in [0,1]$. We can apply the induction hypotheses H1—H3 to G'. In particular, $f\colon x \in [0,e] \mapsto \mathsf{OptCost}_{G'}(q_{\min},x)$ is a cost function satisfying the requirements of item H2. Writing f_1, ..., f_n for the successive affine functions constituting f, we pick the smallest index i such that for every $j > i$, function f_j has slope less than or equal to $-P(q_{\min})$. If $i > 0$, we note $[u,v]$ the domain of f_i (see Fig. 9).

Lemma 7. *If $i=0$, for all $(q,x) \in Q \times [0,e]$, $\mathit{OptCost}_G(q,x) = \mathit{OptCost}_{G'}(q,x)$. If $i>0$, for all $(q,x) \in Q \times [v,e]$, $\mathit{OptCost}_G(q,x) = \mathit{OptCost}_{G'}(q,x)$.*

We now explain how to compute $\mathsf{OptCost}_G(q_{\min},x)$ for $x \in [u,v]$; we prove the following lemma:

Lemma 8. *If $i > 0$, then for all $(q, x) \in Q \times [u, v]$, we have $\mathit{OptCost}_G(q_{\min}, x) = (v - x)P(q_{\min}) + f(v)$.*

The optimal cost in states (q, x) with $x \in [u, v]$ can then be computed by considering the PTG$_f$ G'', restricted to $x \in [u, v]$, and obtained from G' by making $q_{\min}$ a goal location with cost function equal to $x \mapsto \mathsf{OptCost}_G(q_{\min}, x)$, which is then viewed as an outside cost function, see Fig. 9.

We can then repeat the procedure above on the interval $[0, u]$ (*i.e.* by setting $e = u$): compute $f' \colon x \mapsto \mathsf{OptCost}_{G'}(q_{\min}, x)$ with $x \in [0, u]$, select an interval $[u', v']$ where f'_i has slope larger than or equal to $-P(q_{\min})$, and so on, replace that part with an affine function with slope $-P(q_{\min})$, and continue with the interval $[0, u']$. We now explain why this process terminates: since they have slopes strictly greater than $-P(q_{\min})$, f_i and f'_i are fragments of outside cost functions, according to hypothesis H2. If they have different slopes, then they are obviously parts of two different fragments of outside cost functions. If they have the same slopes, then they are fragments of two different parts of outside cost functions, since they are joined by affine functions with slopes less than (or equal to $-P(q_{\min})$). Since there are only finitely many affine functions constituting the outside cost functions, our procedure terminates.

At each step of the procedure above, we can also compute (ε, N)-acceptable strategies, and merge them.

6 Conclusion

In this paper we have proven that optimal cost for arbitrary priced timed games with one clock is a computable problem, and that ε-optimal memoryless strategies may effectively be obtained. The complexity of our procedure is quite high, running in 3-EXPTIME, while the best known lower bound for this problem is PTIME. Our future works of course include tightening these bounds.

As a consequence of our result it may be shown that the iterative semi-algorithm proposed in [10] always terminates for priced timed games with one clock. Cost functions cost^i_G are inductively defined, which for any location $q \in Q$ and any clock value v, give the optimal cost of winning from the state (q, v) within at most i steps (we count the number of steps in a run ρ by the number of delay-and-action fractions). Now Theorem 4 ensures that we can find a fixed N such that for *any* $\varepsilon > 0$ we can compute an (ε, N)-acceptable strategy. In particular this guarantees that we can find ε-optimal strategies which are guaranteed to win within $N \cdot |Q|$ steps for any $\varepsilon > 0$. Consequently, $\langle \mathsf{cost}^i_G \rangle_{i=1}^{\infty}$ (the semi-algorithm of [10]) converges after at most $N \cdot |Q|$ iterations to the optimal cost of winning. A prototype implementation of this iterative algorithm is available at `http://www.cs.aau.dk/~illum/tools/1ptga/`.

As future work we would like to determine what happens with priced timed games using two clocks, but this seems really difficult as our approach heavily relies on the fact that there is only one clock.

References

1. Y. Abdeddaïm, E. Asarin, and O. Maler. Scheduling with timed automata. 2006.
2. R. Alur, M. Bernadsky, and P. Madhusudan. Optimal reachability in weighted timed games. In *Proc. 31st Intl Coll. Automata, Languages and Programming (ICALP'04)*, LNCS 3142, p. 122–133. Springer, 2004.
3. R. Alur, C. Courcoubetis, and Th. A. Henzinger. Computing accumulated delays in real-time systems. In *Proc. 5th Intl Conf. Computer Aided Verification (CAV'93)*, LNCS 697, p. 181–193. Springer, 1993.
4. R. Alur and D. Dill. A theory of timed automata. *Theor. Comp. Science*, 126(2):183–235, 1994.
5. R. Alur, S. La Torre, and G. J. Pappas. Optimal paths in weighted timed automata. In *Proc. 4th Intl Workshop Hybrid Systems: Computation and Control (HSCC'01)*, LNCS 2034, p. 49–62. Springer, 2001.
6. G. Behrmann, A. Fehnker, Th. Hune, K. G. Larsen, P. Pettersson, J. Romijn, and F. Vaandrager. Minimum-cost reachability for priced timed automata. In *Proc. 4th Intl Workshop Hybrid Systems: Computation and Control (HSCC'01)*, LNCS 2034, p. 147–161. Springer, 2001.
7. P. Bouyer, Th. Brihaye, and N. Markey. Improved undecidability results on weighted timed automata. *Inf. Proc. Letters*, (5):188–194, June 2006.
8. P. Bouyer, E. Brinksma, and K. G. Larsen. Staying alive as cheaply as possible. In *Proc. 7th Intl Workshop Hybrid Systems: Computation and Control (HSCC'04)*, LNCS 2993, p. 203–218. Springer, 2004.
9. P. Bouyer, E. Brinksma, and K. G. Larsen. Optimal infinite scheduling for multi-priced timed automata. *Form. Meth. in Syst. Design*, 2006. To appear.
10. P. Bouyer, F. Cassez, E. Fleury, and K. G. Larsen. Optimal strategies in priced timed game automata. In *Proc. 24th Conf. Foundations of Software Technology & Theoretical Computer Science (FST&TCS'04)*, LNCS 3328, p. 148–160. Springer, 2004.
11. Th. Brihaye, V. Bruyère, and J.-F. Raskin. On optimal timed strategies. In *Proc. 3rd Intl Conf. Formal Modeling and Analysis of Timed Systems (FORMATS'05)*, LNCS 3821, p. 49–64. Springer, 2005.
12. S. La Torre, S. Mukhopadhyay, and A. Murano. Optimal-reachability and control for acyclic weighted timed automata. In *Proc. 2nd IFIP Intl Conf. Theoretical Computer Science (IFIPTCS'02)*, IFIP Conf. Proc. 223, p. 485–497. Kluwer, 2002.
13. K. G. Larsen, G. Behrmann, E. Brinksma, A. Fehnker, Th. Hune, P. Pettersson, and J. Romijn. As cheap as possible: Efficient cost-optimal reachability for priced timed automata. In *Proc. 13th Intl Conf. Computer Aided Verification (CAV'01)*, LNCS 2102, p. 493–505. Springer, 2001.
14. K. G. Larsen and J. I. Rassmussen. Optimal conditional reachability for multi-priced timed automata. In *Proc. 8th Intl Conf. Foundations of Software Science and Computation Structures (FoSSaCS'05)*, LNCS 3441, p. 234–249. Springer, 2005.
15. J. I. Rasmussen, K. G. Larsen, and K. Subramani. Resource-optimal scheduling using priced timed automata. In *Proc. 10th Intl Conf. Tools and Algorithms for the Construction and Analysis of Systems (TACAS'04)*, LNCS 2988, p. 220–235. Springer, 2004.
16. UPPAAL CORA. `http://www.cs.aau.dk/~behrmann/cora/`, Jan. 2006.

A Appendix

A.1 Overview of the transformation

Fig. 10 summarizes the transformations of Section 4.

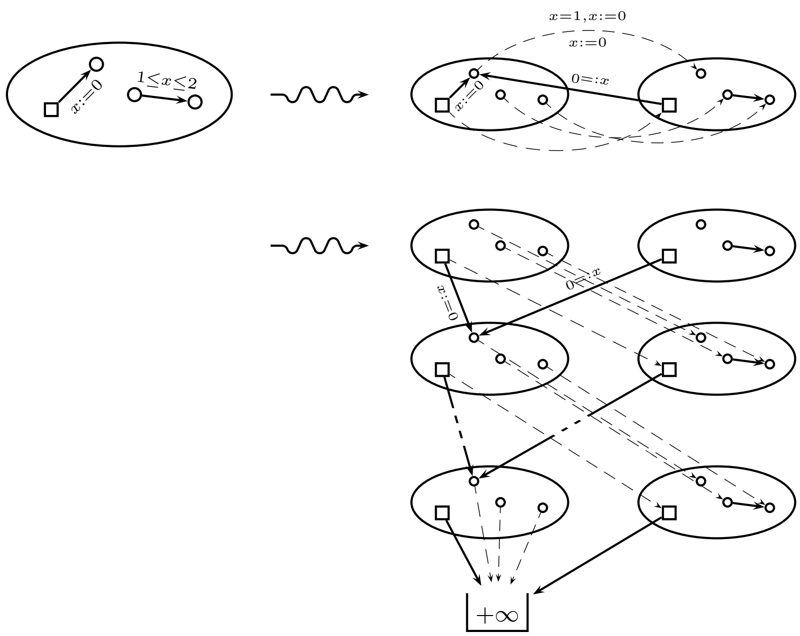

Fig. 10. Overview of the transformation

A.2 Proofs of Section 4 - PTG$_f$ bounded by 1

Proposition 2. *For every state (q, x) in G, $\mathsf{OptCost}_G(q, x) = \mathsf{OptCost}_{G'}(f(q, x))$. Moreover, for every $\varepsilon > 0$ and $N \in \mathbb{N}$, given an (ε, N)-acceptable strategy in G', we can compute an (ε, N)-acceptable strategy in G, and vice-versa.*

Proof. We define the relation $\mathcal{R} \subseteq \mathsf{Run}_G(Q \times \mathbb{R}_+, Q \times \mathbb{R}_+) \times \mathsf{Run}_{G'}(Q' \times \mathbb{R}_+, Q' \times \mathbb{R}_+)$ inductively as follows:

– $(q, x)\mathcal{R}f(q, x)$,
– if we assume $\varrho\mathcal{R}\varrho'$, then:
  - $\varrho\mathcal{R}\Big(\varrho' \to (q_{[\alpha,\alpha+1]}, 0)\Big)$ if $last(\varrho') = (q_{[\alpha-1,\alpha]}, 1)$ for some $0 < \alpha < M$
  - $\varrho \to (last(\varrho){+}t)\mathcal{R}\varrho' \to (last(\varrho'){+}t)$ for every t with no invariant violation
  - $\Big(\varrho \to (q', 0)\Big)\mathcal{R}\Big(\varrho' \to (q'_{[0,1]}, 0)\Big)$ if a resetting transition has been taken

- $\Big(\varrho \to (q', x')\Big)\mathcal{R}\Big(\varrho' \to (q'_{[\alpha,\alpha+1]}, x'-\alpha)\Big)$ if a non-resetting transition has been taken

It is obvious that for every run ϱ in G, there are several ϱ' in G' such that $\varrho\mathcal{R}\varrho'$. Conversely, it is also obvious that to every run ϱ' in G' we can associate a canonical run ϱ with no time-stuttering such that $\varrho\mathcal{R}\varrho'$. Moreover, if $\varrho\mathcal{R}\varrho'$, then $\mathsf{cost}_G(\varrho) = \mathsf{cost}_{G'}(\varrho')$.

Let σ be a strategy in G. Let ϱ' be a run in G', and take ϱ its canonical corresponding run (as described above). We define

$$\sigma'(\varrho') = \begin{cases} q'_{[\alpha,\alpha+1]} & \text{if } \mathit{last}(\varrho') = (q'_{[\alpha-1,\alpha]}, 1) \text{ and } \alpha < M \\ \lambda & \text{otherwise, and if } \sigma(\varrho) = \lambda \\ q'_{[\alpha,\alpha+1]} & \text{otherwise and if } \sigma(\varrho) = q' \end{cases}$$

σ' is a strategy s.t. $\varrho\mathcal{R}\varrho'$ implies $\sigma(\varrho)\mathcal{R}\sigma'(\varrho')$. Obviously, $\mathsf{Cost}_G((q,x),\sigma) = \mathsf{Cost}_{G'}((q_{[\alpha,\alpha+1]}, x'), \sigma')$ if $(q,x)\mathcal{R}(q_{[\alpha,\alpha+1]}, x')$. Moreover, if σ is (ε, N)-acceptable, then so is σ'.

Let σ' be a strategy in G'. Let ϱ be a run in G, and pick one run ϱ' in G' such that $\varrho\mathcal{R}\varrho'$ and such that $\mathit{last}(\varrho') = (q_{[\alpha,\alpha+1]}, x)$ with $\alpha+1 < M$ and x integer, implies $x = 0$. Define $\sigma(\varrho) = \lambda$ iff $\sigma'(\varrho') = \lambda$, and $\sigma(\varrho) = q$ if $\sigma'(\varrho') = q_{[\alpha,\alpha+1]}$. Then it is obvious that $\mathsf{Cost}_G((q,x),\sigma) = \mathsf{Cost}_{G'}((q_\star, x'), \sigma')$ if $(q,x)\mathcal{R}(q_\star, x')$. Moreover, if σ' is (ε', N')-acceptable, then so is σ. □

A.3 Proofs of Section 4 - removing resets

Proposition 3. *For every state (q,x) in the game G, $\mathsf{OptCost}_G((q,x))$ equals $\mathsf{OptCost}_{G'}(((q,0),x))$. Moreover, for every $\varepsilon' > 0$ and $N' \in \mathbb{N}$, given an (ε', N')-acceptable strategy σ' in G', we can compute a $(2\varepsilon', N')$-acceptable strategy σ in G.*

Proof. First, given a (possibly memoryfull) strategy σ' on G', it is easy to build a (memoryfull) strategy σ on G mimicking the same behavior: a run ϱ in G containing $i < n$ resetting transitions can be mapped into a similar run ϱ' in G', and $\sigma(\varrho)$ is defined according to $\sigma'(\varrho')$. The definition is similar if $i = n$ and $\sigma'(\varrho') \neq r$. Otherwise, the strategy might be defined by playing any (fixed) possible move in G. With this definition, if σ' is winning from some state $((q,0),x)$, then σ is winning from (q,x) in G, and $\mathsf{Cost}_{G,\sigma}((q,x)) \leq \mathsf{Cost}_{G',\sigma'}(((q,0),x))$. As a result:

Lemma 9. *For every state (q,x) in G, $\mathsf{OptCost}_G((q,x)) \leq \mathsf{OptCost}_{G'}(((q,0),x))$.*

We now prove the converse inequality, using the same technique: given a strategy σ in G, we define a strategy σ' in G' with lower (or equal) cost. The main idea of the construction is that if a resetting transition is taken twice, then exactly the same state will be crossed twice (because there is only one clock), it is thus not necessary, and we can apply the first time already the strategy that we play the second time.

Assume σ is winning after run ϱ_0. Let t be a resetting transition, and let R be the set of runs $\varrho \in \mathsf{Run}_{G,\sigma}(\varrho_0)$ ending by transition t. There exists some $\varrho_t \in R$ such that no run $\varrho' \in \mathsf{Run}_{G,\sigma}(\varrho_0 \cdot \varrho_t)$ goes through t again (otherwise σ would not be winning after ϱ_0). We then define σ_t after ϱ_0 by

$$\sigma_t(\varrho_0 \cdot \varrho) = \begin{cases} \sigma(\varrho_0 \cdot \varrho) & \text{if } \varrho \text{ does not contain any reset} \\ \sigma(\varrho_0 \cdot \varrho_t \cdot \varrho') & \text{if } \varrho = \varrho'' \cdot \varrho' \text{ with } \varrho'' \text{ ending by } t \end{cases}$$

The strategy σ_t is winning after ϱ_0, obviously $\mathsf{Cost}_G(\sigma_t, \varrho_0) \leq \mathsf{Cost}_G(\sigma, \varrho_0)$. Moreover, all runs $\varrho \in \mathsf{Run}_{G,\sigma_t}(\varrho_0)$ contain at most one occurrence of transition t. We do such a construction for every initial run ϱ_0 such that σ is winning from ϱ_0. We then do the same thing for all resetting transitions, and construct $\sigma_{t,t'}$ from σ_t, and so on. In that way, we define the strategy σ_{L_r} (with L_r a list of all resetting transitions) after ϱ_0 when σ is winning after ϱ_0. If σ is not winning after ϱ, then we set $\sigma_{L_r}(\varrho) = \sigma(\varrho)$.

The strategy σ_{L_r} is winning after ϱ_0 in G if σ was winning after ϱ_0, and $\mathsf{Cost}_G(\sigma_{L_r}, \varrho_0) \leq \mathsf{Cost}_G(\sigma, \varrho_0)$. Moreover, all runs in $\mathsf{Run}_{G,\sigma_{L_r}}(\varrho_0)$ go through each resetting transition at most once. If σ was not winning after ϱ_0, then so is σ_{L_r}.

Now, let ϱ' be a run in G' not ending in r but ending in copy i, and ϱ be its projection in G (defined in the obvious way). If ϱ is compatible with σ_{L_r}, we define $\sigma'(\varrho') = \lambda$ if $\sigma_{L_r}(\varrho) = \lambda$, $\sigma'(\varrho') = (\sigma_{L_r}(\varrho), i')$ if there exists an i' s.t. this corresponds to a valid move in G' (i' is either i if a non-resetting transition is taken, or $i+1$ is a resetting has to be taken), and $\sigma'(\varrho') = r$ otherwise. If ϱ is not compatible with σ_{L_r}, $\sigma'(\varrho')$ is defined as any (fixed) possible next location.

Now, let ϱ' be a run in G' starting in some state $((q,0),x)$, and ϱ be its projection in G. If $\mathsf{Cost}_G(\sigma, \varrho)$ is finite, then so is $\mathsf{Cost}_G(\sigma_{L_r}, \varrho)$. By construction of σ', there is a one-to-one correspondence between the runs compatible with σ' and the runs compatibles with σ_{L_r}. This entails in particular that

$$\mathsf{Cost}_{G'}(\sigma', \varrho') \leq \mathsf{Cost}_G(\sigma_{L_r}, \varrho).$$

As a consequence:

Lemma 10. *For every state (q,x) in G, $\mathsf{OptCost}_{G'}(((q,0),x)) \leq \mathsf{OptCost}_G((q,x))$.*

Combining Lemmas 9 and 10, we have proved the first part of Proposition 3.

Let σ' be an (ε', N')-acceptable strategy in G'. We write σ'_i for the restriction of σ' to states in the i-th copy of G, *i.e.*, $\sigma'_i(q,x) = \sigma'((q,i),x)$. For each winning state $s = (q,x)$ in G, we know that $((q,0),x)$ is winning in G' (because there exists a strategy for which all compatible runs fire each resetting transition at most once). Let $i(s)$ be the maximum number of resetting transitions along the runs compatible with σ' from $((q,0),x)$. That is, $i(s)$ is such that σ' is winning from $((q,n-i(s)),x)$, but not from $((q,n-i(s)+1),x)$. If s is not winning, we let $i(s) = n$. We then define $\sigma(s) = \sigma'_{i(s)}(s)$.

Lemma 11. *σ is $2\varepsilon'$-optimal.*

Proof. Let (q, x) be a state in G, and $i = i((q, x))$. There is nothing to prove if (q, x) is not winning. If it is winning, let $\varrho = (q, x) \to (q', x')$ be a run compatible with σ. It corresponds to a run ϱ' compatible with σ' in G' from $((q, n - i), x)$ in the obvious way, with $\mathsf{cost}(\varrho) = \mathsf{cost}(\varrho')$. That is, $\mathsf{Cost}_G(\sigma, (q, x)) = \mathsf{Cost}_{G'}(\sigma', ((q, n - i), x))$. Now, it cannot be the case that $\mathsf{Cost}_{G'}(\sigma', ((q, n - i), x)) > \mathsf{Cost}_{G'}(\sigma', ((q, 0), x)) + \varepsilon'$, since it would entail that σ' is not ε'-optimal in $((q, n - i), x)$.

Thus

$$\begin{aligned}
\mathsf{Cost}_G(\sigma, (q, x)) &\le \mathsf{Cost}_{G'}(\sigma', ((q, 0), x)) + \varepsilon' \\
&\le \mathsf{OptCost}_{G'}(((q, 0), x)) + 2\varepsilon' \\
&\le \mathsf{OptCost}_G((q, x)) + 2\varepsilon'.
\end{aligned}$$

□

It is clear that σ' fulfills the other requirements needed for being $(2\varepsilon', N')$-acceptable. This concludes the proof of Proposition 3. □

A.4 Proofs of Section 5 – Case $q_{\min}$ controllable

Lemma 5. *For every $(q, x) \in (Q \setminus \{q_{\min}\}) \times [0, 1]$, we have $\mathit{OptCost}_G(q, x) = \mathit{OptCost}_{G'}((q, 0), x)$. For every $x \in [0, 1]$, $\mathit{OptCost}_G(q_{\min}, x) = \mathit{OptCost}_{G'}(q_{\min}, x)$.*

Proof. For writing convenience, we write the proof for $q \neq q_{\min}$, but it is exactly the same for $q_{\min}$.

Obviously, for every (q, x), $\mathsf{OptCost}_G(q, x) \le \mathsf{OptCost}_{G'}((q, 0), x)$. To prove the converse inequality, we first state several lemmas.

Lemma 12. *Let σ be a (possibly memoryfull) strategy in G. Then for every ϱ_0 such that σ is winning from ϱ_0, for every location $q \in Q$, for every $\varrho_1 \in \mathit{Run}_{G,\sigma}(\varrho_0)$,*

- *either for every run $\varrho_2 \in \mathit{Run}_{G,\sigma}(\varrho_0 \cdot \varrho_1)$, q does not appear in ϱ_2,*[5]
- *or there exists $\varrho_2 \in \mathit{Run}_{G,\sigma}(\varrho_0 \cdot \varrho_1)$ ending in q such that for every $\varrho_3 \in \mathit{Run}_{G,\sigma}(\varrho_0 \cdot \varrho_1 \cdot \varrho_2)$, q does not appear in ϱ_3.*

Proof. Assume the contrary, and take witnesses q and ϱ_0. There exists $\varrho_1 \in \mathsf{Run}_{G,\sigma}(\varrho_0)$ such that for every $\varrho_2 \in \mathsf{Run}_{G,\sigma}(\varrho_0 \cdot \varrho_1)$ ending in location q, there exists $\varrho_3 \in \mathsf{Run}_{G,\sigma}(\varrho_0 \cdot \varrho_1 \cdot \varrho_2)$, q appears along ϱ_3. Pick such a $\varrho_1 \in \mathsf{Run}_{G,\sigma}(\varrho_0)$, and take $\varrho_2 \in \mathsf{Run}_{G,\sigma}(\varrho_0 \cdot \varrho_1)$ such that ϱ_2 does contain a location different from q (if this was not possible, it would mean that all runs in $\mathsf{Run}_{G,\sigma}(\varrho_0 \cdot \varrho_1)$ stay forever in q, which contradicts the fact that σ is winning from ϱ_0). Then pick $\varrho_3 \in \mathsf{Run}_{G,\sigma}(\varrho_0 \cdot \varrho_1 \cdot \varrho_2)$ such that q appears along ϱ_3. And so on...

Inductively, we can then build a run ϱ in $\mathsf{Run}_{G,\sigma}(\varrho_0)$ which goes through q infinitely often. This run ϱ is not winning, which contradicts the fact that σ is winning from ϱ_0. □

[5] Note that q can be the first configuration of ϱ_2, but it does not appear afterwards.

Lemma 13. *Let σ be a (possibly memoryfull) strategy in G. Then there exists a strategy σ' such that for every ϱ_0 not containing $q_{\min}$ or ending in $q_{\min}$ such that σ is winning from ϱ_0, σ' is winning from ϱ_0, $\mathsf{Cost}_G(\sigma', \varrho_0) \leq \mathsf{Cost}_G(\sigma, \varrho_0)$, and for every run $\varrho \in \mathsf{Run}_{G,\sigma'}(\varrho_0)$, there is only one occurrence of location $q_{\min}$ along ϱ (possibly with some delay in $q_{\min}$).*

Proof. We define σ' as follows:

- if ϱ has no occurrence of $q_{\min}$, then $\sigma'(\varrho) = \sigma(\varrho)$,
- if ϱ ends in $(q_{\min}, x)$, then define (applying Lemma 12) $\varrho' \in \mathsf{Run}_{G,\sigma}(\varrho)$ ending in $(q_{\min}, y)$ such that for all $\varrho'' \in \mathsf{Run}_{G,\sigma}(\varrho \cdot \varrho')$, ϱ'' does not contain any occurrence of $q_{\min}$. Then for every $0 \leq t < y - x$, we define $\sigma'(\varrho \xrightarrow{t} (q_{\min}, x + t)) = \lambda$, and for every $\sigma'(\varrho \cdot (q_{\min}, x) \xrightarrow{y-x} (q_{\min}, y) \cdot \varrho'') = \sigma(\varrho \cdot \varrho' \cdot \varrho'')$.

Let ϱ_0 be a run not containing $q_{\min}$ or ending in $q_{\min}$ such that σ is winning from ϱ_0. Let $\varrho \in \mathsf{Run}_{G,\sigma'}(\varrho_0)$ be a maximal run. We distinguish two cases:

- either $q_{\min}$ does not appear along ϱ. In that case, $\varrho \in \mathsf{Run}_{G,\sigma}(\varrho_0)$. As σ is winning, ϱ ends in a final state, and $\mathsf{Cost}_G(\varrho) \leq \mathsf{Cost}_G(\sigma, \varrho_0)$.
- or $q_{\min}$ appears along ϱ. In that case we can decompose ϱ as $\varrho_1 \cdot (q_{\min}, x) \xrightarrow{t} (q_{\min}, y) \cdot \varrho_2$ with $\varrho_1 \in \mathsf{Run}_{G,\sigma}(\varrho_0)$ and $\varrho_2 \in \mathsf{Run}_{G,\sigma}(\varrho_0.\varrho_1.\varrho')$ for some ϱ' (with constraints as above). As σ is a winning strategy from ϱ_0, ϱ ends in a final location. Moreover, $\mathsf{Cost}_G(\varrho) = \mathsf{Cost}_G(\varrho_1) + \mathsf{Cost}_G((q_{\min}, x) \xrightarrow{t} (q_{\min}, y)) + \mathsf{Cost}_G(\varrho_2)$. Note that as $q_{\min}$ is a location with smallest cost rate and ϱ' has duration t, we have that $\mathsf{Cost}_G((q_{\min}, x) \xrightarrow{t} (q_{\min}, y)) \leq \mathsf{Cost}_G(\varrho')$. Thus, $\mathsf{Cost}_G(\varrho) \leq \mathsf{Cost}_G(\varrho_1 \cdot \varrho' \cdot \varrho_2) \leq \mathsf{Cost}_G(\sigma, \varrho_0)$.

Thus, σ' is winning from ϱ_0 and $\mathsf{Cost}_G(\sigma', \varrho_0) \leq \mathsf{Cost}_G(\sigma, \varrho_0)$. □

Now, assume that σ is a winning strategy in game G, and construct σ' as in Lemma 13. We define a strategy σ'' in G' by mimicking σ'. It is easy to prove that σ'' is winning from any state $((q, 0), x)$ in G', and that $\mathsf{Cost}_{G'}(\sigma'', \varrho) = \mathsf{Cost}_G(\sigma', \varrho)$, as soon as ϱ does not go through a location $(q, 1)$. In the end, $\mathsf{OptCost}_{G'}((q, 0), x) \leq \mathsf{OptCost}_G(q, x)$, and we have proved Lemma 5. □

Proposition 6. *Strategy σ is winning and there exists a fixed (independant of ε) integer N s.t. σ is (ε, N)-acceptable.*

Proof. We begin with some intermediary results:

Lemma 14. *Let G_0 be a PTG_f. If ϱ is a run in G_0 starting in (q, x) and ending in (q', x'), and compatible with some given memoryless strategy σ, then*

$$\mathsf{Cost}_{G_0}(\sigma, (q, x)) \geq \mathit{cost}(\varrho) + \mathsf{Cost}_{G_0}(\sigma, (q', x')).$$

Moreover, if σ is ε-optimal, then

$$\mathit{cost}(\varrho) \leq \mathit{OptCost}_{G_0}((q, x)) - \mathit{OptCost}_{G_0}((q', x')) + \varepsilon.$$

The first result is rather obvious since the cost of a strategy is the supremum of the costs of all runs compatible with that strategy. The second result follows by ε-optimality of σ.

We then have the following lemma:

Lemma 15. *If for state* (q,x)*, we have*

$$\mathsf{Cost}_{G'}(\sigma',((q,1),x)) \leq \mathit{OptCost}_G(q_{\min},x), \tag{2}$$

and if $\varrho' \in \mathit{Run}_{G',\sigma'}((q,1),x)$ *ends in* $((q',1),x')$*, then*

$$\mathsf{Cost}_{G'}(\sigma',((q',1),x')) \leq \mathit{OptCost}_G(q_{\min},x').$$

Proof. Applying Lemma 14 in game G', we have

$$\mathsf{Cost}_{G'}(\sigma',((q,1),x)) \geq \mathsf{cost}(\varrho') + \mathsf{Cost}_{G'}(\sigma',((q',1),x')).$$

Now, clearly, $\mathsf{cost}(\varrho') \geq P(q_{\min})(x'-x)$, since ϱ' runs in locations of the SCC, thus with costs larger than $P(q_{\min})$. From the hypothesis, we get

$$\mathsf{Cost}_{G'}(\sigma',((q',1),x')) \leq \mathsf{OptCost}_G(q_{\min},x) - P(q_{\min})(x'-x).$$

Now, one possible strategy in G from $(q_{\min},x)$ is to delay until $(q_{\min},x')$. The optimal cost is thus less than or equal to the cost of this strategy, that is,

$$\mathsf{OptCost}_G(q_{\min},x) \leq P(q_{\min})(x'-x) + \mathsf{OptCost}_G(q_{\min},x').$$

Thus,

$$\mathsf{Cost}_{G'}(\sigma',((q',1),x')) \leq \mathsf{OptCost}_G(q_{\min},x').$$

□

Lemma 16. *If Equation* (2) *holds for some state* (q,x)*, then*

$$\mathsf{Cost}_{G'}(\sigma',((q,1),x)) - \varepsilon' \leq \mathsf{Cost}_{G'}(\sigma',((q,0),x)) \leq \mathsf{Cost}_{G'}(\sigma',((q,1),x)) + \varepsilon'.$$

Proof. As $\mathsf{OptCost}_{G'}((q,0),x) \leq \mathsf{OptCost}_{G'}((q,1),x)$ and σ' is ε'-optimal, we have

$$\mathsf{Cost}_{G'}(\sigma',((q,0),x)) \leq \mathsf{Cost}_{G'}(\sigma',((q,1),x)) + \varepsilon.$$

Now, assume

$$\mathsf{Cost}_{G'}(\sigma',((q,0),x)) < \mathsf{Cost}_{G'}(\sigma',((q,1),x)) - \varepsilon'. \tag{3}$$

From (2), we get

$$\mathsf{Cost}_{G'}(\sigma',((q,0),x)) < \mathsf{OptCost}_G(q_{\min},x) - \varepsilon'. \tag{4}$$

We prove that $q_{\min}$ does not appear in any of the runs compatible with σ' from $((q,0),x)$. Otherwise, pick $\varrho' \in \mathsf{Run}_{G',\sigma'}((q,0),x)$ ending in some $((q_{\min},0),x')$. Applying Lemma 14,

$$\begin{aligned}\mathsf{Cost}_{G'}(\sigma',((q,0),x)) &\geq \mathsf{cost}(\varrho') + \mathsf{Cost}_{G'}(\sigma',(q_{\min},x'))\\ &\geq P(q_{\min})(x'-x) + \mathsf{OptCost}_{G'}(q_{\min},x')\\ &\geq P(q_{\min})(x'-x) + \mathsf{OptCost}_{G}(q_{\min},x')\\ &\geq \mathsf{OptCost}_{G}(q_{\min},x).\end{aligned}$$

This contradicts Equation (4), and thus $q_{\min}$ does not appear in any of the runs compatible with σ' from $((q,0),x)$. This entails that we can mimic σ' from $((q,1),x)$. Thus, $\mathsf{OptCost}_{G'}((q,1),x) \leq \mathsf{Cost}_{G'}(\sigma',((q,0),x))$, and as σ' is ε'-optimal, $\mathsf{Cost}_{G'}(\sigma',((q,1),x)) \leq \mathsf{OptCost}_{G'}((q,1),x)+\varepsilon'$, which in turn contradicts Equation (3), and establishes our result. □

We now complete the proof of Proposition 6. Let $\varrho \in \mathsf{Run}_{G,\sigma}(q,x)$. If all along ϱ, condition (2) fails to hold, then σ always mimics σ', and ϱ corresponds to a run compatible with σ' from $((q,0),x)$, and thus has almost optimal cost (up to ε').

Otherwise, from some state (q',x') on (possibly (q,x)), σ will mimic σ' from $((q',1),x')$. Let c be the cost of the portion of ϱ between (q,x) and (q',x'):

$$(q,x) \xrightarrow[c]{\sigma' \text{ in component } 0} (q',x') \xrightarrow{\sigma' \text{ in component } 1} \text{winning state}$$

Then,

$$\begin{aligned}\mathsf{cost}(\varrho) &\leq c + \mathsf{Cost}_G(\sigma,(q',x'))\\ &\leq c + \mathsf{Cost}_{G'}(\sigma',((q',1),x'))\\ &\leq c + \mathsf{Cost}_{G'}(\sigma',((q',0),x')) + \varepsilon'\\ &\leq c + \mathsf{OptCost}(q',x') + 2\varepsilon'.\end{aligned}$$

The second last inequality is from Lemma 16, the last one comes from ε'-optimality of σ' and the fact that (q',x') and $((q',0),x')$ have the same optimal costs.

By Lemma 14, we also have

$$\mathsf{Cost}_{G'}(\sigma',((q,0),x)) \geq c + \mathsf{Cost}_{G'}(\sigma',((q',0),x')).$$

Now, by ε'-optimality of σ', we have

$$\begin{aligned}\mathsf{Cost}_{G'}(\sigma',((q,0),x)) &\leq \mathsf{OptCost}_G(q,x)+\varepsilon'\\ \mathsf{OptCost}_G(q',x') &\leq \mathsf{Cost}_{G'}(\sigma',((q',0),x'))\end{aligned}$$

which yields

$$c \leq \mathsf{OptCost}_G(q,x) - \mathsf{OptCost}_G(q',x') + \varepsilon'.$$

Thus

$$\mathsf{cost}(\varrho) \le \mathsf{OptCost}_G(q,x) + 3\varepsilon' = \mathsf{OptCost}_G(q,x) + \varepsilon.$$

This inequality holds for every run ϱ compatible with σ from a winning state (q,x). Thus σ is winning and ε-optimal.

Let N_0 be the number of intervals defining the piecewise-affine function $x \mapsto \mathsf{OptCost}_{G'}(q_{\min},x)$. There is thus a partition with $N'.N_0$ intervals such that the functions $x \mapsto \mathsf{OptCost}_{G'}(q_{\min},x)$ and for every location q, $x \mapsto \mathsf{Cost}_{G'}(\sigma',(q,x))$ are affine on every interval. Thus in such a given interval, for every location q, there is at most one smallest α such that $\mathsf{OptCost}_{G'}(q_{\min},\alpha) = \mathsf{Cost}_{G'}(\sigma',((q,1),\alpha))$, and we can prove that the strategy of component 0 is played before α whereas it is the strategy of component 1 which is played after α. Thus, if we fix $N = |Q|.N'.N_0$, then strategy σ is (ε,N)-acceptable, which concludes the proof of property H3. □

A.5 Proofs of Section 5 – Case $q_{\min}$ uncontrollable

Lemma 7. *If $i = 0$, for all $(q,x) \in Q\times[0,e]$, $\mathit{OptCost}_G(q,x) = \mathit{OptCost}_{G'}(q,x)$. If $i > 0$, for all $(q,x) \in Q \times [v,e]$, $\mathit{OptCost}_G(q,x) = \mathit{OptCost}_{G'}(q,x)$.*

Proof. We pick some $\varepsilon > 0$, and prove that $\mathsf{OptCost}_G(q,x) \le \mathsf{OptCost}_{G'}(q,x)+\varepsilon$. To that aim, we use hypothesis H3: we pick N' as defined by hypothesis H3, we let $\varepsilon' = \varepsilon/(N'+1)$, and take an (ε',N')-acceptable strategy σ' in G' on the interval $[0,e]$. Let ϱ be a run compatible with σ' in G starting from some state (q,x). We have:

$$\begin{aligned}\varrho : (q,x) \ &\rightsquigarrow^* (q_{\min},x_1) \xrightarrow{c_1} (q_{\min},x_1+d_1)\\ &\cdots\\ &\rightsquigarrow^* (q_{\min},x_n) \xrightarrow{c_n} (q_{\min},x_n+d_n) \rightsquigarrow^* \ldots\end{aligned}$$

where $c_i = P(q_{\min})\cdot d_i$ and $\rightsquigarrow^*$ denotes a sequence where $q_{\min}$ is not visited. If the number of returns in $q_{\min}$ is larger than N', then there exists an interval $[c,d]$ over which the strategy σ' is constant, and during which ϱ goes twice in $q_{\min}$. This indicates that, in G', the opponent can play in such a way that the game goes back to $q_{\min}$ infinitely often over $[c,d]$, so that σ' is not winning in state (q,x) in G'. In that case, we obviously have $\mathsf{OptCost}_G(q,x) \le \mathsf{OptCost}_{G'}(q,x)$.

We now assume that ϱ visits $q_{\min}$ at most N' times. Applying Lemma 14 on the $(i+1)$-st $\rightsquigarrow^*$-part (denoted by ϱ_{i+1}, and seen here as an outcome of σ' in G'), we have

$$\mathsf{cost}(\varrho_{i+1}) \le \mathsf{OptCost}_{G'}(q_{\min},x_i+d_i) - \mathsf{OptCost}_{G'}(q_{\min},x_{i+1}) + \varepsilon'. \tag{5}$$

Similarly for the first part of the trajectory:

$$\mathsf{cost}(\varrho_{i+1}) \le \mathsf{Cost}_{G'}(\sigma',(q,x)) - \mathsf{OptCost}_{G'}(q_{\min},x_1).$$

Thus:

$$\begin{aligned}\mathsf{cost}(\varrho) \le (\mathsf{Cost}_{G'}(\sigma',(q,x)) - f(x_1)) + c_1 + \\ + (f(x_1+d_1) - f(x_2) + \varepsilon') + c_2 \\ \dots \\ + (f(x_{n-1}+d_{n-1}) - f(x_n) + \varepsilon') + c_n + f(x_n+d_n) \\ \le \mathsf{Cost}_{G'}(\sigma',(q,x)) + N'\varepsilon'\end{aligned}$$

where the last inequality follows from the fact that the slopes of f are all less than or equal to $P(q_{\min})$, so that we have $f(x_i) - f(x_i+d_i) \ge P(q_{\min})d_i$.

We have thus proved that, for every $x \in [0,e]$,

$$\begin{aligned}\mathsf{OptCost}_G(q,x) &\le \mathsf{Cost}_G(\sigma',(q,x)) + N'\varepsilon' \\ &\le \mathsf{OptCost}_{G'}(q,x) + (N'+1)\varepsilon' \\ &\le \mathsf{OptCost}_{G'}(q,x) + \varepsilon.\end{aligned}$$

For the case $i > 0$, the proof above still applies on the interval $[v,e]$, where $[v,e]$ is the definition interval of the affine functions f_j for $j > i$. $\square$

Lemma 8. *If $i > 0$, then for all $(q,x) \in Q \times [u,v]$, we have $\mathit{OptCost}_G(q_{\min},x) = (v-x)P(q_{\min}) + f(v)$.*

Proof. First, as waiting in $q_{\min}$ until v is one possible strategy for the opponent, we have

$$\mathsf{OptCost}_G(q_{\min},x) \ge (v-x)P(q_{\min}) + \mathsf{OptCost}_G(q_{\min},v) = (v-x)P(q_{\min}) + f(v).$$

We prove that the converse inequality holds, by picking an arbitrary $\varepsilon > 0$ and showing that

$$\mathsf{OptCost}_G(q_{\min},x) \le (v-x)P(q_{\min}) + f(v) + \varepsilon.$$

Again, we pick N' satisfying the conditions of item H3, let $\varepsilon' = \varepsilon/(N'+1)$, and pick a strategy σ' given by hypothesis H3. Let $x \in [u,v]$, and ϱ be an outcome of σ', played in G, from $(q_{\min},x)$. As previously, ϱ can be depicted as:

$$\begin{aligned}\varrho : (q_{\min},x) \rightsquigarrow^* (q_{\min},x_1) \xrightarrow{c_1} (q_{\min},x_1+d_1) \\ \dots \\ \rightsquigarrow^* (q_{\min},x_n) \xrightarrow{c_n} (q_{\min},x_n+d_n) \rightsquigarrow^* \dots\end{aligned}$$

with possibly $x = x_1$. We assume that $x_n+d_n \in [u,v]$, and that after state $(q_{\min},x_n+d_n)$, ϱ never goes back to $q_{\min}$ within $[u,v]$. Again, we may assume that ϱ visits at most N' times location $q_{\min}$. Equation (5) still applies, and since the slope of f is now larger than $-P(q_{\min})$, it yields

$$\begin{aligned}\mathsf{cost}(\varrho_{i+1}) &\le \mathsf{OptCost}_{G'}(q_{\min},x_i+d_i) - \mathsf{OptCost}_{G'}(q_{\min},x_{i+1}) + \varepsilon' \\ &\le P(q_{\min}) \cdot (x_{i+1} - x_i - d_i) + \varepsilon'.\end{aligned}$$

We end up with

$$\mathsf{cost}(\varrho) \le P(q_{\min}) \cdot (x_n + d_n - x) + f(x_n + d_n) + N'\varepsilon'.$$

Again, from the inequality on the slope of f in $[u, v]$, we get

$$f(x_n + d_n) \le f(v) + P(q_{\min})(v - x_n - d_n),$$

from which we get

$$\mathsf{cost}(\varrho) \le P(q_{\min})(v - x) + f(v) + N'\varepsilon'.$$

This holds for any outcome of σ' in G from $(q_{\min}, x)$ with $x \in [u, v]$. This proves the inequality we claimed:

$$\mathsf{OptCost}_G(q_{\min}, x) \le \mathsf{Cost}_G(\sigma', (q_{\min}, x)) \le P(q_{\min})(v - x) + f(v) + \varepsilon.$$

□

A.6 Complexity of our procedure

We evaluate here the complexity of our algorithm. We only consider the case of $[0,1]$-PTG$_f$ based on an SCC (or having only one single location) and without resets.

We begin with evaluating the number of segments of affine functions that appear in the optimal cost functions (of all locations). Assume the SCC contains n locations, and f outside cost functions made of a total of p segments of affine functions. We write $N(n, p)$ for the total number of affine segments that appear in the optimal cost functions of all the locations of that SCC (as we will see, it does not depend on f). We then have $N(1, p) \le 2p$ (because in the worst case, one part of each segment of the outside cost functions appear in the optimal cost, and it might be needed to delay in the location between each segment), and the following recursive equations:

$$N(n, p) \le N(n-1, p + N(1, p + N(n-1, p+1))) \qquad \text{(case } q_{\min} \text{ controllable)}$$
$$N(n, p) \le p \cdot (1 + N(n-1, p+1)) \qquad \text{(case } q_{\min} \text{ uncontrollable)}$$

The inequalities follow from the constructions presented above: for instance, in the case where $q_{\min}$ is controllable, the optimal costs are obtained by first computing the optimal costs in the rightmost copy of the game of Fig. 5 (this yields $N(n-1, p+1)$ segments), then add those cost functions as outside cost functions to $q_{\min}$, compute the optimal cost function of $q_{\min}$ (which is bounded by $N(1, p + N(n-1, p+1))$), add this cost function as an outside cost function of the leftmost copy, and compute the optimal costs in that game.

It can easily be checked then that $N(n, p) \le 2^{2^{n \cdot 2^n}} \cdot (p+n)^{2^{2^n}}$, i.e., $N(n, p)$ is at most triply exponential. The time needed to compute optimal costs and almost-optimal strategies is then readily shown to be in 3-EXPTIME. This holds for one single SCC, but also extends to combinations of SCCs, and thus to the full class of PTG$_f$.