

A Faster Algorithm for Solving One-Clock Priced Timed Games*

Thomas Dueholm Hansen[†] Rasmus Ibsen-Jensen[†] Peter Bro Miltersen[†]

Abstract

One-clock priced timed games is a class of two-player, zero-sum, continuous-time games that was defined and thoroughly studied in previous works. We show that one-clock priced timed games can be solved in time $m12^n n^{O(1)}$, where n is the number of states and m is the number of actions. The best previously known time bound for solving one-clock priced timed games was $2^{O(n^2+m)}$, due to Rutkowski. For our improvement, we introduce and study a new algorithm for solving one-clock priced timed games, based on the sweep-line technique from computational geometry and the strategy iteration paradigm from the algorithmic theory of Markov decision processes. As a corollary, we also improve the analysis of previous algorithms due to Bouyer, Cassez, Fleury, and Larsen; and Alur, Bernadsky, and Madhusudan.

1 Introduction

Priced timed *automata* and priced timed *games* are classes of one-player and two-player zero-sum real-time games played on finite graphs that were defined and thoroughly studied in previous works [2, 4, 3, 16, 1, 7, 9, 6, 8, 12, 14]. Synthesizing (near-)optimal strategies for priced timed games has many practical applications in embedded systems design; we refer to the cited papers for references.

Informally (for formal definitions, see the sections below), a priced timed game is played by two players on a finite directed labeled multi-graph. The vertices of the graph are called *states*, with some states belonging to Player 1 (or the Minimizer) and the other states belonging to Player 2 (or the Maximizer). We shall denote by n the total number of states of the game under consideration and m the total number of arcs (actions). Player 1 is trying to play the game to termination as cheaply as possible, while Player 2 is trying to make Player 1 pay as dearly as possible for playing. At any point in time, some particular state is the *current* state. The player controlling the state decides when to leave the current state and which arc to follow when doing so. For each arc, there is an associated *cost*. Each state has an associated *rate* of expense per time unit associated with waiting in the state. The above setup is further refined by the introduction of a finite number of *clocks* that can informally be thought of as “stop watches”. In particular, some arcs may have associated a *reset* event for a clock. If the corresponding transition is taken, that clock is reset to 0. Also, an arc may have an associated clock and time interval. When the arm of the clock is in the interval, the corresponding transition can be taken; otherwise it can not. With three or more clocks, the problem of solving priced timed games is known not to be computable [6]. In this paper, we focus on the computable case of solving one-clock priced timed games. We shall refer to these as PTGs. We shall furthermore single out an important, particularly clean, special case of PTGs. We shall refer to this class as *simple priced timed games*, SPTGs. In an SPTG, time runs from 0 to 1, the single clock is never reset, and there are no restrictions on when transitions may be taken. A slightly more general class of games was called “[0,1]-PTGs without resets” by Bouyer et al. [8].

*Work supported by the Sino-Danish Center for the Theory of Interactive Computation, funded by the Danish National Research Foundation and the National Science Foundation of China (under the grant 61061130540). The authors acknowledge support from the Center for research in the Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. Thomas Dueholm Hansen is a recipient of the Google Europe Fellowship in Game Theory, and this research is supported in part by this Google Fellowship.

[†]Department of Computer Science, Aarhus University, Denmark. E-mail: {tdh, rij, pbmiltersen}@cs.au.dk.

As is the case in general for two-player zero-sum games, informally, a priced timed game is said to have a *value* v if Player 1 and Player 2 are both able to guarantee (or approximate arbitrarily well) a total cost of v when the game is played. The guarantees are obtained when players commit to (near-)optimal *strategies* when playing the game. Player 1, who is trying to minimize cost, may (approximately) guarantee the value from above, while Player 2, who is trying to maximize cost, may (approximately) guarantee the value from below. Clearly, in general, the value of a one-clock priced timed game will be a function $v(q, t)$ of the initial state q and the initial setting t of the single clock. Bouyer *et al.* [8] showed that the value $v(q, t)$ exists¹ and that for any state q , the value function $t \rightarrow v(q, t)$ is a piecewise linear function of t . By *solving* a game, we mean computing an explicit description of all these functions (i.e., lists of their line segments). From such an object, near-optimal strategies can be synthesized.

Figure 1 shows an SPTG with $n = 5$ states. Circles are controlled by Player 1 and squares are controlled by Player 2. States and actions have been annotated with rates and costs. If no cost is given for an action it has cost zero. The figure also includes graphs of the value functions. Actions are shown in black and gray, and an optimal strategy profile is shown along the x -axis of the value functions by using these colors – more precisely, it is the optimal strategy found by our algorithm. Waiting is shown as white.

If both players follow the indicated optimal strategies, then the play that starts with state 3 as the current state at time 0, is as follows:

1. At state 3 at time 0, Player 1 waits until time $\frac{1}{3}$ and then changes the current state to state 2.
2. At state 2 at time $\frac{1}{3}$, Player 2 waits until time $\frac{2}{3}$ and then changes the current state to state 4.
3. At state 4 at time $\frac{2}{3}$, Player 1 does not wait, but immediately changes the current state to state 3.
4. At state 3 at time $\frac{2}{3}$, Player 1 waits until time 1 and then changes the current state to state 1.
5. At state 1 at time 1, Player 2 can not wait, and immediately changes the current state to state \perp , a special state indicating that play has terminated.

Notice that the play waits in state 3 twice. This may seem like a counter-intuitive property of a play where the players play optimally. In fact, the game can be generalized to a family, such that the game with n states has a state that is visited $O(n)$ times in some optimal play.

The contributions of this paper are the following.

1. *A polynomial time Turing-reduction from the problem of solving general PTGs to the problem of solving SPTGs.* The best previous result along these lines was a Turing-reduction from the general case to the case of “[0,1]-PTGs without resets” by Bouyer *et al.* [8]. Our reduction is a polynomial time reduction reducing solving a general PTG to solving at most $(n+1)(2m+1)$ SPTGs, while the previous reduction is an exponential time reduction.
2. *A novel algorithm for solving SPTGs, based on very different techniques than previously used to solve PTGs.* In particular, our algorithm is based on applications of a technique from computational geometry: the *sweep-line* technique of Shamos and Hoey [15], applied to the linear arrangement resulting when the graphs of all value functions are superimposed in a certain way. Also, an extension of Dijkstra’s algorithm due to Khachiyan *et al.* [13] is a component of the algorithm. We believe that an implementation of this algorithm and the reduction could provide an attractive alternative to the current state-of-the-art tools for solving PTGs or various special cases (e.g., such as those of UPPAAL, <http://uppaal.org> or HyTech <http://embedded.eecs.berkeley.edu/research/hytech/>), which all seem to be based on a value-iteration based algorithm independently devised by Bouyer, Cassez, Fleury, and Larsen [7]; and Alur, Bernadsky, and Madhusudan [1]. We shall refer to that algorithm as the BCFL-ABM algorithm.

¹Players in general cannot guarantee the value exactly, but only approximate it arbitrarily well – one of the particular appealing aspects of SPTGs is that they *do* have exactly optimal strategies! This is in contrast to both the general case and [0,1]-PTGs without resets.

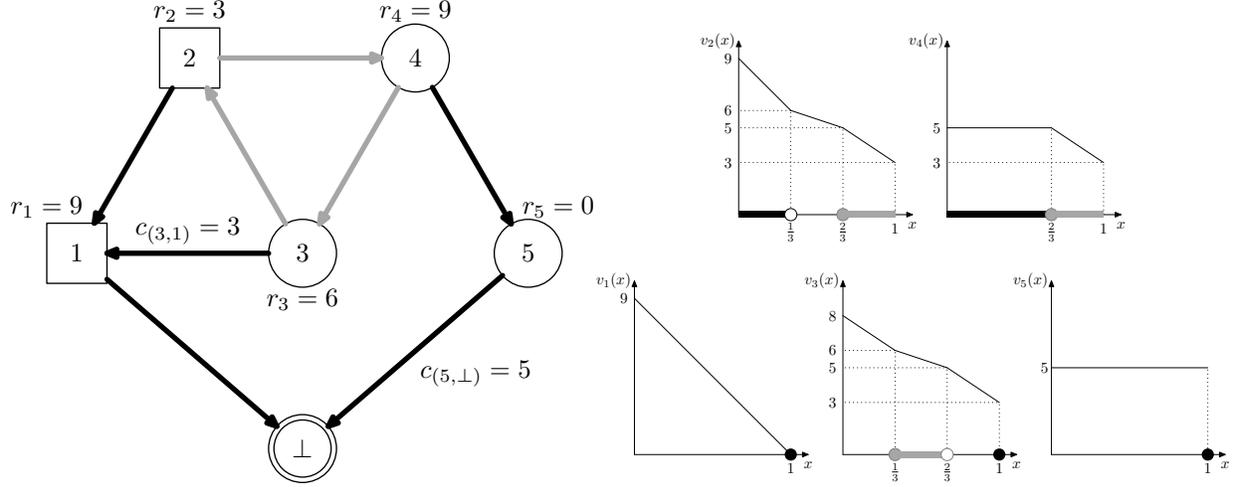


Figure 1: Example of an SPTG, showing value functions and an optimal strategy profile.

3. *A worst case analysis of our algorithm as well as an improved worst case analysis of the BCFL-ABM algorithm.* Interestingly, the analysis of the algorithms is quite indirect: We analyze a different algorithm for a subproblem (priced games, see section 2), namely the *strategy iteration* algorithm, also used to solve Markov decision processes and various other classes of two-player zero-sum games played on graphs, and relate the analysis of this algorithm to our algorithm. To summarize the result of the analysis, it is convenient to introduce the parameter $L = L(G)$ of an SPTG to be the total number of distinct time coordinates of left endpoints of the linear segments of all value functions of G . Note that the parameter L is very natural, as L is a lower bound on the size of the explicit description of these value functions, i.e., the output of the algorithms under consideration. We show:

- (a) For an SPTG G , we have that $L(G) \leq \min\{12^n, \prod_{k \in S} (|A_k| + 1)\}$, where S is the set of states and A_k the set of actions in state k . The best previous bound on $L(G)$ was $2^{O(n^2)}$, due to Rutkowski [14].
- (b) The worst case time complexity of our new algorithm is $O((m + n \log n)L)$. In particular, the algorithm combined with the reduction solves general PTGs in time $m12^n n^{O(1)}$. The best previous worst case bound for any algorithm solving PTGs was $2^{O(n^2+m)}$, due to Rutkowski [14], who gave this bound for an alternative algorithm, due to him.
- (c) The worst case number of iterations of the BCFL-ABM algorithm is $\min\{12^n, \prod_{k \in S} (|A_k| + 1)\}m \cdot n^{O(1)}$ for general PTGs, significantly improving an analysis of Rutkowski. (An "iteration" is a natural unit of time, specific to the algorithm – each iteration may take considerable time, as entire graphs of value functions are manipulated during an iteration).
- (d) For the special case of PTGs with all rates being 1 (i.e., all states are equally expensive to wait in) and all transition costs being 0 (i.e., Player 1 wants to minimize the time used), our algorithm combined with the reduction runs in time $O(nm(\min(m, n^2) + n \log n))$. The previously best algorithm for solving this special case (called *timed reachability games*) is an exponential time algorithm due to Jurdzinski and Trivedi [12].
- (e) For one-clock priced timed automata (the special case of priced timed games, where all states belong to Player 1), our algorithm combined with the reduction runs in time $O(mn^3(\min(m, n^2) + n \log n))$. This seems to be the best worst case bound known for solving these.

The above bounds hold if we assume a unit-cost Real RAM model of computation, which is a natural model of computation for the algorithms considered (that previous analyses also seem to have implicitly assumed).

The algorithms can also be analyzed in Boolean models of computation (such as the log cost integer RAM), as a rational valued input yields a rational valued output. Bounding the bit length of the numbers computed by straightforward inductive techniques, we find that this no more than squares the above worst case complexity bounds. The somewhat tedious analysis establishing this is not included in this version of the paper.

1.1 History of problem and related research

Priced timed automata (or weighted timed automata) were first introduced by Alur, Torre, and Pappas [3] and Behrmann *et al.* [4]. They showed that priced timed automata (viewed as one-player games) can be solved in exponential time. Even before the introduction of priced timed automata, a special case was studied by Alur and Dill [2]. They show this case to be PSPACE-hard even for automata where all states have rate 1 and all actions cost 0. Bouyer, Brihaye, Bruyere, and Raskin [5] showed that the problem of solving priced timed automata is in PSPACE. I.e., the problem is PSPACE-complete when there is no limit on the number of clocks.

Bouyer, Cassez, Fleury, and Larsen [7] and Alur, Bernadsky, and Madhusudan [1] independently introduced the notion of priced timed games and also both considered value iteration algorithms for solving priced timed games. Finding the value of a priced timed game with many clocks is a hard problem. Even with only 3 clocks, finding the value becomes undecidable for priced timed games, as shown by Bouyer, Brihaye and Markey [6]. They improved a similar result of Brihaye, Bruyere, and Raskin [9] for 5 clocks. Hence, various special cases have been studied. For timed reachability games, Jurdzinski and Trivedi [12] showed the decision problem to be in EXP and EXP-complete for 2 or more clocks.

For the case with only one clock the problem becomes computable, as shown by Brihaye, Bruyere, and Raskin [9]. Bouyer, Larsen, Markey, and Rasmussen [8] gave an explicit triple exponential time bound on the complexity of solving this problem. This was further improved to $2^{O(n^2+m)}$ by Rutkowski [14].

1.2 Organization of paper

Our algorithm is most naturally presented in three stages, adding more complications to the model at each stage. First, in section 2, we show how the strategy iteration paradigm can be used to solve *priced games*, where the temporal aspects of the games are not present. In section 3, we show how the algorithm extends to simple priced timed games. Finally, in section 4, we show how solving the general case of one-clock priced-timed games can be reduced to the case of simple priced timed games in polynomial time.

In terms of the list of contributions above, contribution 1) is Lemma 4.8. The algorithm of contribution 2) is `SolveSPTG` of Figure 5. Contribution 3a) is Theorem 3.10, contribution 3b) is Theorem 3.11, contribution 3c) is Theorem 4.10, contribution 3d) is Theorem 4.11 and contribution 3e) is Theorem 4.12.

2 Priced games

In this section, we introduce *priced games*. To accommodate lexicographic utilities which will be necessary for subsequent sections, we shall consider priced games with utilities in domains other than \mathbb{R} . In this section, we fix any ordered Abelian group $(\mathfrak{R}, +, -, 0, \leq)$ for the set of possible utilities. We let $\mathfrak{R}_{\geq 0}$ be the set of non-negative elements in \mathfrak{R} . In subsequent sections, we will either have $\mathfrak{R} = \mathbb{R}$ or $\mathfrak{R} = \mathbb{R} \times \mathbb{R}$ with lexicographic order. In the latter case, we write (x, y) as $x + y\epsilon$, where we informally think of ϵ as an infinitesimal. In addition to utilities in the group \mathfrak{R} , we also allow the utility ∞ (modeling non-termination).

A *priced game* G is given by a finite set of *states* $S = [n] = \{1, \dots, n\}$, a finite set of *actions* $A = [m] = \{1, \dots, m\}$. The set S is partitioned into S_1 and S_2 , with S_i being the set of states belonging to Player i . Player 1 is also referred to as the *minimizer* and Player 2 is referred to as the *maximizer*. The set A is partitioned into $(A_k)_{k \in S}$, with A_k being the set of actions available in state k . Furthermore, define $A^i = \bigcup_{k \in S_i} A_k$. Each action $j \in A$ has an associated non-negative *cost* $c_j \in \mathfrak{R}_{\geq 0} \cup \{\infty\}$ and an associated *destination* $d(j) \in S \cup \{\perp\}$, where \perp is a special *terminal state*. Note that G can be interpreted as a directed

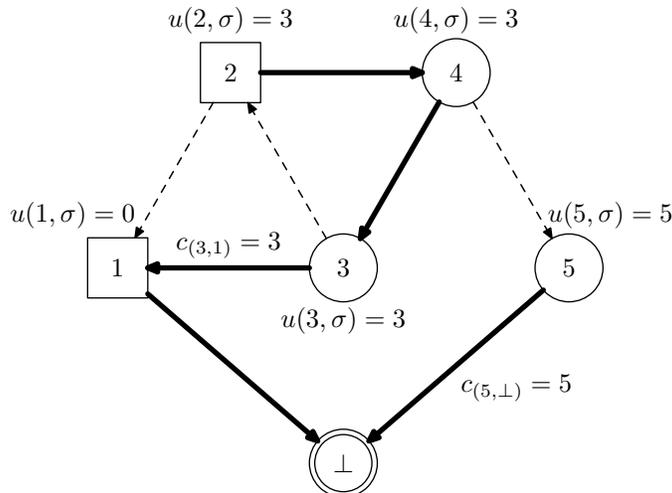


Figure 2: Example of a priced game and a strategy profile σ .

weighted graph. In fact, a priced game can be viewed as a single source shortest path problem from the point of view of Player 1, with the exception that an adversary, Player 2, controls some of the decisions.

A *positional strategy* for Player i is a map σ_i of S_i to A , with $\sigma(k) \in A_k$ for each $k \in S$. A pair of strategies (or *strategy profile*) $\sigma = (\sigma_1, \sigma_2)$ defines a maximal path $P_{k_0, \sigma} = (k_0, k_1, \dots)$, from each $k_0 \in S \cup \{\perp\}$, possibly ending at \perp , such that $d(\sigma(k_i)) = k_{i+1}$ for all $i \geq 0$. Note that σ can be naturally interpreted as a map from S to A . Let $\ell(k, \sigma)$ be the length of $P_{k, \sigma}$. The path $P_{k, \sigma}$ defines a *payoff* $u(k, \sigma) \in \mathbb{R} \cup \{\infty\}$, paid by Player 1 to Player 2, as follows:

$$u(k, \sigma) = \begin{cases} \infty & \text{if } \ell(k, \sigma) = \infty \\ 0 & \text{if } k = \perp \\ c_{\sigma(k)} + u(d(\sigma(k)), \sigma) & \text{otherwise} \end{cases}$$

I.e., the payoff is the total cost of the path $P_{k, \sigma}$ from k to the terminal state \perp , or ∞ if $P_{k, \sigma}$ does not reach \perp .

The *lower value* $\underline{v}(k)$ of a state k is defined by $\underline{v}(k) = \max_{\sigma_2} \min_{\sigma_1} u(k, \sigma_1, \sigma_2)$. A strategy σ_2 is called *optimal*, if for all states k , we have $\sigma_2 \in \operatorname{argmax}_{\sigma_2} \min_{\sigma_1} u(k, \sigma_1, \sigma_2)$. Similarly, the *upper value* $\bar{v}(k)$ of a state k is defined by $\bar{v}(k) = \min_{\sigma_1} \max_{\sigma_2} u(k, \sigma_1, \sigma_2)$ and a strategy σ_1 is called optimal if for all k , $\sigma_1 \in \operatorname{argmin}_{\sigma_1} \max_{\sigma_2} u(k, \sigma_1, \sigma_2)$. Khachiyan *et al.* [13] observed that $\underline{v}(k) = \bar{v}(k)$, i.e., that priced games have *values* $v(k) := \underline{v}(k) = \bar{v}(k)$. They also showed how to find these values and optimal strategies efficiently using a variant of Dijkstra's algorithm. The **ExtendedDijkstra** algorithm is shown in Figure 3, with v being the vector of values. Viewing a priced game as a single source shortest path problem, it is not surprising that it can be solved by a Dijkstra-like algorithm. Intuitively, if an arc to be taken by Player 2 would be optimal for Player 1, Player 2 will, if possible, do anything else and, informally, "delete" the arc.

Figure 2 shows an example of a priced game. The round vertices are controlled by Player 1, the minimizer, and the square vertices are controlled by Player 2, the maximizer. Bold arrows indicate actions used by a strategy profile σ , and dashed arrows indicate unused actions. Actions are labeled by their cost, except if the cost is zero. Finally, the states have been annotated by the values. Note that σ is an optimal strategy profile.

We say that σ_1 is a *best response* to σ_2 if $\sigma_1 \in \operatorname{argmin}_{\sigma_1} u(k, \sigma_1, \sigma_2)$, for all $k \in S$. Similarly, σ_2 is a *best response* to σ_1 if $\sigma_2 \in \operatorname{argmax}_{\sigma_2} u(k, \sigma_1, \sigma_2)$, for all $k \in S$. A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is a *Nash equilibrium* if σ_1 is a best response to σ_2 , and σ_2 is a best response to σ_1 . The following is a standard lemma that establishes the connection between Nash equilibria and values of zero-sum games.

Lemma 2.1 *If $\sigma = (\sigma_1, \sigma_2)$ is a Nash equilibrium, then $v(k) = u(k, \sigma)$ for all $k \in S$.*

Function ExtendedDijkstra(G)

```

( $v(\perp), v(1), \dots, v(n)$ )  $\leftarrow$  ( $0, \infty, \dots, \infty$ );
while  $S \neq \emptyset$  do
  ( $k, j$ )  $\leftarrow$   $\underset{k \in S, j \in A_k}{\operatorname{argmin}}$   $c_j + v(d(j))$ ;
  if  $k \in S_1$  or  $|A_k| = 1$  then
     $v(k) \leftarrow c_j + v(d(j))$ ;
     $\sigma(k) \leftarrow j$ ;
     $S \leftarrow S \setminus \{k\}$ ;
  else
     $A_k \leftarrow A_k \setminus \{j\}$ ;
return ( $v, \sigma$ );

```

Figure 3: The ExtendedDijkstra algorithm of Khachiyan *et al.* [13] for solving priced games.

Proof: Assume that either σ_1 or σ_2 is not optimal. We will show that (σ_1, σ_2) is not a Nash equilibrium for play starting in some state of the game. Assume, without loss of generality, that σ_1 does not guarantee Player 1 the payoff $v(k)$ for play starting at k . There are two cases.

- Case 1: $u(k, \sigma_1, \sigma_2) \leq v(k)$. In this case, Player 2 can deviate from σ_2 to play a best response to σ_1 at state k . Since σ_1 does, by assumption, not guarantee Player 1 $v(k)$, this will yield a larger payoff than $v(k)$, i.e., the deviation improves payoff for Player 2 and (σ_1, σ_2) is therefore not a Nash equilibrium for play starting at k .
- Case 2: $u(k, \sigma_1, \sigma_2) > v(k)$. In this case, Player 1 can deviate to play an optimal strategy σ_1^* . By definition of optimal, this improves his payoff to $v(k)$ and (σ_1, σ_2) is therefore not a Nash equilibrium for play starting at k .

□

We shall present a different algorithm for solving priced games, following the general *strategy iteration* pattern [11]. This algorithm will be extended to priced timed games in the next sections. Let σ be a strategy profile. For each state $k \in S$, we define the *valuation* $\nu(k, \sigma) = (u(k, \sigma), \ell(k, \sigma))$. I.e., the valuation of a state k for strategy profile σ is the payoff for k combined with the length of the path $P_{k, \sigma}$. If $\nu(k, \sigma) = (\infty, \infty)$ we write $\nu(k, \sigma) = \infty$. We say that an action $j \in A_k$ from state k is an *improving switch for Player 1* if:

$$(c_j + u(d(j), \sigma), 1 + \ell(d(j), \sigma)) < \nu(k, \sigma)$$

Where we order pairs lexicographically, with the first component being most significant. I.e., an improving switch for Player 1 either produces a path from k of smaller cost or with the same cost and smaller length. Similarly, $j \in A_k$ is an improving switch for Player 2 if:

$$(c_j + u(d(j), \sigma), 1 + \ell(d(j), \sigma)) > \nu(k, \sigma)$$

Lemma 2.2 *Let $\sigma = (\sigma_1, \sigma_2)$ be a strategy profile such that for both players i , there are no improving switches in A^i . Then σ_1 and σ_2 are optimal.*

Proof: By Lemma 2.1 it is enough to show that (σ_1, σ_2) is a Nash equilibrium for play starting in each state of the game.

Let σ'_1 be a best response to σ_2 , and let $\sigma' = (\sigma'_1, \sigma_2)$. Assume, for the sake of contradiction, that there exists a $k_0 \in S$ such that $u(k_0, \sigma') < u(k_0, \sigma)$. Let k_i be the i 'th state on the path $P_{k_0, \sigma'}$. I.e., $k_{i+1} = d(\sigma'(k_i))$.

Either $P_{k_0, \sigma'}$ leads to the terminal state \perp , or $P_{k_0, \sigma'}$ is an infinite path ending in a cycle. The second case is impossible since that would imply $u(k_0, \sigma') = \infty < u(k_0, \sigma)$.

Since σ'_1 is a best response to σ_2 , we have $u(k_i, \sigma') \leq u(k_i, \sigma)$ for all i . Also, since $u(\perp, \sigma') = u(\perp, \sigma) = 0$ and $u(k_0, \sigma') < u(k_0, \sigma)$, when $P_{k_0, \sigma'}$ leads to the terminal state, there must exist an index i such that $u(k_{i+1}, \sigma') = u(k_{i+1}, \sigma)$ and $u(k_i, \sigma') < u(k_i, \sigma)$. Thus, $\sigma'(k_i)$ is an improving switch for Player 1, and since $\sigma(k_i) \neq \sigma'(k_i)$ we have $k_i \in S_1$; a contradiction.

The argument for Player 2 is similar. Let σ'_2 be a best response to σ_1 , and $\sigma' = (\sigma_1, \sigma'_2)$. Assume that $u(k_0, \sigma') > u(k_0, \sigma)$ for some $k_0 \in S$, and let k_i be the i 'th state along the path $P_{k_0, \sigma'}$. For the case when $P_{k_0, \sigma'}$ leads to the terminal state, the argument is the same except with $<$ and $>$ interchanged.

When $P_{k_0, \sigma'}$ is an infinite path ending in a cycle we must have $u(k_0, \sigma') = \infty > u(k_0, \sigma)$. I.e., $u(k_i, \sigma)$ is finite for all i . Recall that $c_j \geq 0$ for all $j \in A$. For all $k_i \in S_1$, $\sigma(k_i) = \sigma'(k_i)$, and, hence:

$$\nu(k_i, \sigma) = (c_{\sigma(k_i)} + u(k_{i+1}, \sigma), 1 + \ell(k_{i+1}, \sigma)) > \nu(k_{i+1}, \sigma).$$

On the other hand, for all $k_i \in S_2$, $\sigma'(k_i)$ is not an improving switch for Player 2, and, hence:

$$\nu(k_i, \sigma) \geq (c_{\sigma'(k_i)} + u(k_{i+1}, \sigma), 1 + \ell(k_{i+1}, \sigma)) > \nu(k_{i+1}, \sigma).$$

Thus, $P_{k_0, \sigma'}$ leads to a cycle, for which the valuations for σ decrease with each step; a contradiction. \square

Let $B \subseteq A$ be a set of actions such that $|B \cap A_k| \leq 1$ for all $k \in S$, and, for $B \cap A_k \neq \emptyset$, let $j(k, B)$ be the unique action in $B \cap A_k$. Let σ be a strategy profile, and let $\sigma[B]$ be defined as:

$$\sigma[B](k) := \begin{cases} j(k, B) & \text{if } B \cap A_k \neq \emptyset \\ \sigma(k) & \text{otherwise.} \end{cases}$$

If $B = \{j\}$ we also write $\sigma[j]$. If $j \in A$ is not an improving switch for one player, we say that j is *weakly improving* for the other player. We say that $B \subseteq A$ is an *improving set for Player i* if there exists an improving switch $j \in B$ for Player i , and for all $j \in B$, j is weakly improving for Player i .

Lemma 2.3 *Let $\sigma = (\sigma_1, \sigma_2)$ be a strategy profile, and let $B \subseteq A$ be an improving set for Player 1. Then $\nu(k, \sigma[B]) \leq \nu(k, \sigma)$ for all $k \in S$, with strict inequality if $\sigma[B](k_0)$ is an improving switch for Player 1 w.r.t. σ . Similarly, if B is an improving set for Player 2, then $\nu(k, \sigma[B]) \geq \nu(k, \sigma)$ for all $k \in S$, with strict inequality if $\sigma[B](k_0)$ is an improving switch for Player 2 w.r.t. σ .*

Proof: First consider the case where B is an improving set for Player 1. Let $k_0 \in S$. We must show that $\nu(k_0, \sigma[B]) \leq \nu(k_0, \sigma)$ with strict inequality if $\sigma[B](k_0)$ is an improving switch for Player 1 w.r.t. σ . This is clearly true if $\nu(k_0, \sigma) = \infty$. Thus, assume that $\nu(k_0, \sigma) < \infty$.

Let k_i be the i 'th state on the path $P_{k_0, \sigma[B]}$. Since $\sigma[B](k_i)$ is weakly improving for Player 1 we have, for all i :

$$(c_{\sigma[B](k_i)} + u(k_{i+1}, \sigma), 1 + \ell(k_{i+1}, \sigma)) \leq \nu(k_i, \sigma) \tag{1}$$

with strict inequality exactly when $\sigma[B](k_i)$ is an improving switch for Player 1 w.r.t. σ .

From (1), and the fact that $c_j \geq 0$ for all $j \in A$, we get that:

$$\begin{aligned} \nu(k_i, \sigma) &\geq (c_{\sigma[B](k_i)} + u(k_{i+1}, \sigma), 1 + \ell(k_{i+1}, \sigma)) \\ &> (u(k_{i+1}, \sigma), \ell(k_{i+1}, \sigma)) \\ &= \nu(k_{i+1}, \sigma). \end{aligned}$$

Hence, $P_{k_0, \sigma[B]}$ does not lead to a cycle, since the valuations in σ can not strictly decrease along the entire cycle.

Function StrategyIteration(G)

```

while  $\exists$  improving set  $B_1 \subseteq A^1$  for Player 1 w.r.t.  $\sigma$  do
   $\sigma \leftarrow \sigma[B_1]$ ;
  while  $\exists$  improving set  $B_2 \subseteq A^2$  for Player 2 w.r.t.  $\sigma$  do
     $\sigma \leftarrow \sigma[B_2]$ ;
return  $(u(\sigma), \sigma)$ ;

```

Figure 4: The **StrategyIteration** algorithm for solving priced games.

We next show, using backwards induction on i , that $\nu(k_i, \sigma[B]) \leq \nu(k_i, \sigma)$. For the base case, $k_i = \perp$, the statement is clearly true. Otherwise, for $i < \ell(k_0, \sigma[B])$, we get from (1) and the induction hypothesis that:

$$\begin{aligned}
\nu(k_i, \sigma) &\geq (c_{\sigma[B](k_i)} + u(k_{i+1}, \sigma), 1 + \ell(k_{i+1}, \sigma)) \\
&\geq (c_{\sigma[B](k_i)} + u(k_{i+1}, \sigma[B]), 1 + \ell(k_{i+1}, \sigma[B])) \\
&= \nu(k_i, \sigma[B]).
\end{aligned}$$

Note that if $j \in A_{k_i} \cap B$ is an improving switch for Player 1, the first inequality is strict.

The proof for the second case, where B is an improving set for Player 2, is similar. Let $k_0 \in S$. We show that $\nu(k_0, \sigma[B]) \geq \nu(k_0, \sigma)$ with strict inequality if $\sigma[B](k_0)$ is an improving switch for Player 2 w.r.t. σ . Now, this is clearly true if $\nu(k_0, \sigma[B]) = \infty$. If $\nu(k_0, \sigma[B]) < \infty$, it immediately follows that $P_{k_0, \sigma[B]}$ is of finite length. The rest of the proof is identical, but with $<$ and $>$ interchanged. \square

Lemmas 2.2 and 2.3 allow us to define the **StrategyIteration** algorithm as shown in Figure 4. $u(\sigma)$ is the vector of payoffs for σ . The algorithm is a local search algorithm, and Lemma 2.2 ensures that a local optimum is also a global optimum. Player 1 repeatedly performs improving switches while Player 2 always plays a best response to the current strategy of Player 1.

Theorem 2.4 *The StrategyIteration algorithm correctly computes an optimal strategy profile.*

Proof: It immediately follows from Lemma 2.2 that if the **StrategyIteration** algorithm terminates, it correctly computes an optimal strategy profile. Indeed, in order to escape both while-loops neither player i can have an improving switch in A^i w.r.t. σ .

Let $\sigma = (\sigma_1, \sigma_2)$ be the current strategy profile at the beginning of the outer while-loop, and let $\sigma[B_1] = (\sigma'_1, \sigma_2)$. From Lemma 2.3 we know that with each iteration of the inner while-loop the valuations are non-decreasing, with at least one state strictly increasing its valuation. Since there are only finitely many strategies, it follows that the inner while-loop always terminate. Let the resulting strategy profile be $\sigma' = (\sigma'_1, \sigma'_2)$. Then σ'_2 is optimal for the game where Player 1 is restricted to play according to σ'_1 . I.e., σ'_2 is a best response to σ'_1 .

After the first iteration σ_2 is a best response to σ_1 . Then all actions in A^2 are weakly improving for Player 1 w.r.t. σ , and $B = B_1 \cup \{j \in A^2 \mid \exists k \in S_2 : j = \sigma'_2(k) \neq \sigma_2(k)\}$ is an improving set for Player 1. Since $\sigma' = \sigma[B]$ it follows that the valuations are non-increasing and strictly decreasing for at least one state from σ to σ' . Again, since there are only finitely many strategies the outer while-loop is guaranteed to terminate. \square

3 Simple priced timed games

A *simple priced timed game* (SPTG) G is given by a priced game $G' = (S_1, S_2, (A_k)_{k \in S}, (c_j)_{j \in A}, d)$, where $S = S_1 \cup S_2$ and $A = \bigcup_{k \in S} A_k$, and for each state $i \in S$, an associated *rate* $r_i \in \mathbb{R}_{\geq 0}$. We assume that

$A_k \neq \emptyset$ for all $k \in S$.

A SPTG G is played as follows. A pebble is placed on some starting state k_0 and the clock is set to its starting time x_0 . The pebble is then moved from state to state by the players. The current configuration of the game is described by a state and a time, forming a pair $(k, x) \in S \times [0, 1]$.

Assume that after t steps the pebble is on state $k_t \in S_i$, controlled by Player i , at time x_t , corresponding to the configuration (k_t, x_t) . Player i now chooses the next action $j_t \in A_{k_t}$. Furthermore, the player also chooses a delay $\delta_t \geq 0$ such that $x_{t+1} = x_t + \delta_t \leq 1$. The pebble is moved to $d(j_t) = k_{t+1}$. The next configuration is then (k_{t+1}, x_{t+1}) . We write

$$(k_t, x_t) \xrightarrow{j_t, \delta_t} (k_{t+1}, x_{t+1}).$$

The game ends if $k_{t+1} = \perp$.

A *play* of the game is a sequence of steps starting from some configuration (k_0, x_0) . Let

$$\rho = (k_0, x_0) \xrightarrow{j_0, \delta_0} (k_1, x_1) \xrightarrow{j_1, \delta_1} \dots \xrightarrow{j_{t-1}, \delta_{t-1}} (k_t, x_t)$$

be a finite play such that $k_t = \perp$. The outcome of the game, paid by Player 1 to Player 2, is then given by:

$$\text{cost}(\rho) = \sum_{\ell=0}^{t-1} (\delta_\ell r_{k_\ell} + c_{j_\ell}).$$

I.e., for each unit of time spent waiting at a state k Player 1 pays the rate r_k to Player 2. Furthermore, every time an action j is used, Player 1 pays the cost c_j to Player 2. If ρ is an infinite play the outcome is ∞ , and we write $\text{cost}(\rho) = \infty$.

A (positional) strategy for Player i is a map $\pi_i : S_i \times [0, 1] \rightarrow A \cup \{\lambda\}$, where λ is a special delay action. For every $k \in S_i$ and $x \in [0, 1)$, if $\pi_i(k, x) = \lambda$ then we require that there exists a $\delta > 0$ such that for all $0 \leq \epsilon < \delta$, $\pi_i(k, x + \epsilon) = \lambda$. Let $\delta_{\pi_i}(k, x) = \inf\{x' - x \mid x \leq x' \leq 1, \pi_i(k, x') \neq \lambda\}$ be the delay before the pebble is moved when starting in state k at time x for some strategy π_i .

Player i is said to play according to π_i if, when the pebble is in state $k \in S_i$ at time $x \in [0, 1]$, he waits until time $x' = x + \delta_{\pi_i}(k, x)$ and then moves according to $\pi_i(k, x')$. A *strategy profile* $\pi = (\pi_1, \pi_2)$ is a pair of strategies, one for each player. Let Π_i be the set of strategies for Player i , and let Π be the set of all strategy profiles. A strategy profile π is again interpreted as a map $\pi : S \times [0, 1] \rightarrow A \cup \{\lambda\}$. Furthermore, we use $\pi(x)$ to refer to the decisions at a fixed time. I.e., $\pi(x) : S \rightarrow A \cup \{\lambda\}$ is the map defined by $(\pi(x))(k) = \pi(k, x)$.

Let $\rho_{k,x}^\pi$ be the play starting from configuration (k, x) where the players play according to π . Define the *value function* for a strategy profile $\pi = (\pi_1, \pi_2)$ and state k as: $v_k^{\pi_1, \pi_2}(x) = \text{cost}(\rho_{k,x}^\pi)$. For fixed strategies π_1 and π_2 for Player 1 and 2, define the *best response value functions* for Player 2 and 1, respectively, for a state k as:

$$\begin{aligned} v_k^{\pi_1}(x) &= \sup_{\pi_2 \in \Pi_2} v_k^{\pi_1, \pi_2}(x) \\ v_k^{\pi_2}(x) &= \inf_{\pi_1 \in \Pi_1} v_k^{\pi_1, \pi_2}(x) \end{aligned}$$

We again define *lower* and *upper value functions*:

$$\begin{aligned} \underline{v}_k(x) &= \sup_{\pi_2 \in \Pi_2} v_k^{\pi_2}(x) = \sup_{\pi_2 \in \Pi_2} \inf_{\pi_1 \in \Pi_1} v_k^{\pi_1, \pi_2}(x). \\ \bar{v}_k(x) &= \inf_{\pi_1 \in \Pi_1} v_k^{\pi_1}(x) = \inf_{\pi_1 \in \Pi_1} \sup_{\pi_2 \in \Pi_2} v_k^{\pi_1, \pi_2}(x). \end{aligned}$$

Note that inf and sup are used because there are infinitely many strategies. Bouyer *et al.* [8] showed that $\underline{v}_k(x) = \bar{v}_k(x)$. In fact, this was shown for the more general class of *priced timed games* (PTGs) studied in Section 4. Thus, every SPTG has a *value function* $v_k(x) := \underline{v}_k(x) = \bar{v}_k(x)$ for each state k .

A strategy $\pi_i \in \Pi_i$ is *optimal from time x* for Player i if:

$$\forall k \in S, x' \in [x, 1]: v_k^{\pi_i}(x') = v_k(x').$$

Strategies are called *optimal* if they are optimal from time 0. Similarly, a strategy π_i is a *best response* to another strategy π_{-i} from time x if:

$$\forall k \in S, x' \in [x, 1]: v_k^{\pi_i, \pi_{-i}}(x') = v_k^{\pi_{-i}}(x').$$

A strategy profile (π_1, π_2) is called a *Nash equilibrium from time x* if π_1 is a best response to π_2 from time x , and π_2 is a best response to π_1 from time x . As in the case of Lemma 2.1 for priced games, any equilibrium payoff of an SPTG is the value of the game. The exact statement is shown in Lemma 3.1. Since the argument is standard, and similar to the proof of Lemma 2.1, it has been omitted. Just note that instead of considering best responses, which we have not yet showed exist for SPTGs, it suffices to use some better strategy.

Lemma 3.1 *If there exists a strategy profile (π_1, π_2) that is a Nash equilibrium from time x , then $v_k(x') = v_k^{\pi_1, \pi_2}(x')$ for all $k \in S$ and $x' \in [x, 1]$.*

The existence of optimal strategies and best replies is non-trivial. We are, however, later going to prove the following theorem, which, in particular, implies that inf and sup can be replaced by min and max in the definitions of value functions. (The second half of the theorem also holds for general PTGs and was first established by Bouyer *et al.* [8] who furthermore showed that the first half fails for general PTGs.)

Theorem 3.2 *For any SPTG there exists an optimal strategy profile. Also, the value functions are continuous piecewise linear functions.*

Our proof will be algorithmic. Specifically, the algorithm `SolveSPTG` computes a value function of the desired kind. Furthermore, the proof of correctness of `SolveSPTG` (the proof of Theorem 3.11) also yields the existence of exactly optimal strategies.

We refer to the non-differentiable points of the value functions of G as *event points* of G . The number of distinct event points of G is an important parameter in the complexity of our algorithm for solving SPTGs. We denote by $L(G)$ the total number of event points, excluding $x = 1$.

Remark 3.3 Let us remark that strategies are commonly defined as maps from states and times to delays and actions. For instance, $\tau_i : S_i \times [0, 1] \rightarrow [0, 1] \times A$. This is more general than our definition of strategies, since $\tau_i(k, x) = (\delta, a)$ with $\delta > 0$ does not imply that for all $x' \in (x, x + \delta]$ we have $\tau_i(k, x') = (x + \delta - x', a)$, whereas this implication holds for the strategies we use. We choose to use the specialized definition of strategies because it offers a better intuition for understanding the proposed algorithm. It is easy to see that the players can not achieve better values by using the more general strategies. Indeed, let τ_i be some strategy where $\tau_i(k, x) = (\delta, a)$ and $\tau_i(k, x') = (\delta', a')$, such that $[x, x + \delta] \cap [x', x' + \delta'] \neq \emptyset$. Then one of the following two modifications will not make τ_i achieve worse values: $\tau_i(k, x) = (x' + \delta' - x, a')$ or $\tau_i(k, x') = (x + \delta - x', a)$.

3.1 Solving SPTGs

In order to solve an SPTG we make use of a technique similar to the *sweep-line* technique from computational geometry of Shamos and Hoey [15]. Informally, we construct the value functions by moving a sweep-line backwards from time 1 to time 0, and at each time computing the current values based on the later values. The approach is also similar to a technique known in game theory as *backward induction*. The parameter of the induction, the time, is a continuous parameter, however. The BCFL-ABM algorithm also applies a backward induction, but there, the parameter of induction is the number of transitions taken, i.e., a discrete parameter, leading to a *value iteration* algorithm. The formal development of the algorithm follows.

If π is a strategy profile that is optimal from time x , we use π to construct a new strategy profile π' that is optimal from time $x' < x$. More precisely, for $\epsilon > 0$ sufficiently small, we show that there exists a fixed

optimal action for all states for both players for all point of time in the interval $[x', x)$, where $x' = x - \epsilon$. The new profile π' is then obtained from π by using these actions. If the value at time x is known, and the strategies do not change in the interval $[x - \epsilon, x)$, then $v_k(x - \epsilon) = v_k(x) + \epsilon r_k$ if the players wait at state k . The optimal actions can then be found by solving a priced game where waiting is associated with the resulting cost.

Definition 3.4 For a given SPTG $G = (S_1, S_2, (A_k)_{k \in S}, c, d, r)$ and a time $x \in (0, 1]$, let the priced game $G^{x,y} = (S_1, S_2, (A'_k)_{k \in S}, c^x, d')$ be defined by:

$$\begin{aligned} \forall k \in S : \quad & A'_k = A_k \cup \{\lambda_k\} \\ \forall j \in A'_k : \quad & c_j^{x,y} = \begin{cases} v_k(x) + yr_k & \text{if } j = \lambda_k \\ c_j & \text{otherwise} \end{cases} \\ \forall j \in A'_k : \quad & d'(j) = \begin{cases} \perp & \text{if } j = \lambda_k \\ d(j) & \text{otherwise} \end{cases} \end{aligned}$$

We will usually let y be the infinitesimal ϵ , in which case we will simply denote $G^{x,\epsilon}$ by G^x and $c^{x,\epsilon}$ by c^x .

The additional actions λ_k , for $k \in S$, should be interpreted as the additional option of waiting in the SPTG G . Note that the only difference between G^x and $G^{x'}$, for $x \neq x'$, is the costs of the added actions λ_k . Hence, we may interpret a strategy profile σ for G^x as a strategy profile for $G^{x'}$. Also note that G^x is identical to the priced game G' defining G , except that for each state k there is an additional action λ_k corresponding to waiting in that state in the SPTG. Slightly abusing notation, we will interpret actions chosen by σ as also being actions $\pi(x)$ for G , and the actions $\pi(x)$ as forming a strategy profile for G^x .

We will also sometimes write $u(k, \sigma, G^x)$ instead of $u(k, \sigma)$ to clarify which priced game G^x we consider. Since ϵ is an infinitesimal, the payoffs of a strategy profile σ have two components, and we let $u(k, \sigma, G^x) = a(k, \sigma, G^x) + \epsilon b(k, \sigma, G^x)$. Note that $a(k, \sigma, G^x) = u(k, \sigma, G^{x,0})$. Furthermore, for every $x \in (0, 1]$, let $v^x(k) = a^x(k) + \epsilon b^x(k)$ be the value of state k in G^x , and let $\sigma^x = (\sigma_1^x, \sigma_2^x)$ be an optimal strategy profile.

Let $x \in (0, 1]$, let σ be a strategy profile for G^x , let k_0 be a state, and assume that $u(k_0, \sigma, G^x) = a(k_0, \sigma, G^x) + \epsilon b(k_0, \sigma, G^x) < \infty$. We then define $r(k_0, \sigma) = b(k_0, \sigma, G^x)$ to be the rate of the waiting state reached from k_0 when players play according to σ . More precisely, we have $P_{k_0, \sigma} = (k_0, k_1, \dots, k_t)$ where $k_t = \perp$, and it will either be the case that $\sigma(k_{t-1}) = \lambda_{k_{t-1}}$, or that the last action taken was part of the original game. The only actions whose costs have an infinitesimal component are the additional actions λ_k , for $k \in S$. In particular, such infinitesimal costs exactly correspond to the rates of states in G . Hence, in the first case we have $r(k_0, \sigma) = r_{k_{t-1}}$, and in the second case we have $r(k_0, \sigma) = 0$. In both cases $r(k_0, \sigma)$ can be interpreted as the actual rate that will be paid for waiting in the original game. I.e., if we reach \perp before time 1 the rate of waiting there is 0.

Lemma 3.5 Let π be a strategy profile for G that is optimal from time x , and let $x' < x$. If $\pi(x'') = \sigma^x$ for all $x'' \in [x', x)$, then $v_k^\pi(x') = v_k(x) + (x - x')b^x(k)$ for all $k \in S$.

Proof: Let $\rho_{k,x'}^\pi = (k_0, x_0) \xrightarrow{j_0, \delta_0} (k_1, x_1) \xrightarrow{j_1, \delta_1} \dots$, and let t be the maximum index such that $x_t < x$. Since $\pi(x'') = \sigma^x$ for all $x'' \in [x', x)$, we have $\delta_\ell = 0$ for all $\ell < t$ and $\delta_t \geq x - x'$. By splitting the cost of $\rho_{k,x'}^\pi$ into cost accumulated before and after time x , we get:

$$\begin{aligned} v_k^\pi(x') &= \text{cost}(\rho_{k,x'}^\pi) \\ &= \left((x - x')r_{k_t} + \sum_{\ell=0}^{t-1} c_{j_\ell} \right) + v_{k_t}(x) \\ &= a(k, \sigma^x, G^x) + (x - x')b(k, \sigma^x, G^x) \\ &= a^x(k) + (x - x')b^x(k) . \end{aligned}$$

It remains to show that $a^x(k) = v_k(x)$. Recall that $a^x(k) = a(k, \sigma^x, G^x) = u(k, \sigma^x, G^{x,0})$. Observe that since σ^x is optimal for G^x it must also be optimal for $G^{x,0}$. Indeed, if a better value can be achieved in $G^{x,0}$ then, regardless of the infinitesimal component of the payoff achieved by σ^x , it will also be better for G^x . Furthermore, the value of a state k in $G^{x,0}$ must be consistent with $v_k(x)$. It follows that $u(k, \sigma^x, G^{x,0}) = v_k(x)$. □

Recall that σ is optimal for G^x when there are no improving switches w.r.t. σ . Hence, if $0 < y \leq 1$ is the maximum number for which there are no improving switches w.r.t. σ^x in $G^{x,y'}$, for all $y' \in (0, y]$, then σ^x is optimal for all such $G^{x,y'}$. In fact, we will see that $x' = x - y$ is the next event point preceding x . For every action $j \in A$ and time $x \in (0, 1]$, define the function:

$$f_{j,x}(x'') = c_j + a^x(d(j)) + b^x(d(j))(x - x'').$$

Note that if $f_{j,x}(x'') < f_{\sigma^x(k),x}(x'')$, for $k \in S_1$ and $j \in A_k$, then j is an improving switch with respect to σ^x in $G^{x,x-x''}$ for Player 1. Define $\text{NextEventPoint}(G^x)$ as:

$$\max \{0\} \cup \{x' \in [0, x) \mid \exists k \in S, j \in A_k : f_{j,x}(x) \neq f_{\pi(k),x}(x) \wedge f_{j,x}(x') = f_{\pi(k),x}(x')\}.$$

Note that $\text{NextEventPoint}(G^x)$ is well-defined, since there is only one function $f_{j,x}$ for each action $j \in A$.

Lemma 3.6 *Let $x' = \text{NextEventPoint}(G^x)$, then σ^x is optimal for $G^{x,y}$, for all $y \in (0, x - x']$.*

Proof: As discussed prior to stating the lemma, σ^x is optimal for $G^{x,y}$ if neither player i has an improving switch $j \in A^i$ w.r.t. σ^x . Per definition there are no improving switches when y is sufficiently small. Recall that the valuation $\nu(k, \sigma^x, G^{x,y}) = (u(k, \sigma^x, G^{x,y}), \ell(k, \sigma^x, G^{x,y}))$ consists of two components. Although the payoffs $u(k, \sigma^x, G^{x,y})$ change for different y , $\ell(k, \sigma^x, G^{x,y})$ remains the same. Thus, we only need to consider for which payoffs there is an improving switch w.r.t. σ^x .

For $G^{x,y}$ we have, for all $k \in S$:

$$u(k, \sigma^x, G^{x,y}) = c_{\sigma^x(k)}^{x,y} + u(d(\sigma^x(k)), \sigma^x, G^{x,y})$$

Then $j \in A_k$, where $k \in S_1$, is an improving switch for Player 1 w.r.t. σ^x if and only if:

$$c_j + u(d(j), \sigma^x, G^{x,y}) < u(k, \sigma^x, G^{x,y}) = c_{\sigma^x(k)}^{x,y} + u(d(\sigma^x(k)), \sigma^x, G^{x,y}) \iff$$

$$f_{j,x}(x - y) < f_{\sigma^x(k),x}(x - y)$$

The same holds for Player 2 with reversed inequalities. The maximum $x' < x$ for which there is an improving switch j w.r.t. σ^x in $G^{x,y}$ appears when the lines defined by $f_{j,x}$ and $f_{\sigma^x(k),x}$ intersect. If $f_{j,x} = f_{\sigma^x(k),x}$, j never becomes an improving switch, however. $\text{NextEventPoint}(G^x)$ exactly equals such an intersection point, and possibly 0 if there is none. Hence, σ^x is an optimal strategy for $G^{x,y}$ for all $y \in (0, x - x']$, where $x' = \text{NextEventPoint}(G^x)$. □

Lemma 3.7 *Let $x' = \text{NextEventPoint}(G^x)$, and let $\pi = (\pi_1, \pi_2)$ be a strategy profile that is optimal from time x . Then the strategy profile $\pi' = (\pi'_1, \pi'_2)$, defined by:*

$$\pi'(k, x'') = \begin{cases} \sigma^x(k) & \text{if } x'' \in [x', x) \\ \pi(k, x'') & \text{otherwise} \end{cases}$$

is optimal from time x' , and $v_k(x'') = v_k(x) + b^x(k)(x - x'')$, for $x'' \in [x', x)$ and $k \in S$.

Proof: Let us first note that for any strategy profile π'' , the outcome $v_{k_0}^{\pi''}(x_0)$, for some starting configuration (k_0, x_0) , only depends on the choices made by π'' in the interval $[x_0, 1]$. Hence, since π' is the same as π in the interval $[x, 1]$, π' is also optimal from time x .

Let us also note that $v_k(x'') = \infty$ for some $k \in S$ and $x'' \in [0, 1]$ if and only if $v_k(x') = \infty$ for all $x'' \in [0, 1]$, and $v^x(k) = \infty$ for all G^x . Indeed, the value is infinite exactly when the play has infinite length, and this property is independent of time. Hence, costs and rates are of no importance. $v_k(x'')$ is, thus, correctly set to ∞ if $v_k(x) = \infty$. Also, $\sigma^x(k)$ achieves the correct value for k . For the remainder of the proof we focus on the case where $v_k(x) < \infty$. It follows immediately from Lemma 3.5 that the value function has the correct form in the interval $[x', x)$. I.e., $v_k^{\pi'}(x'') = v_k(x) + b^x(k)(x - x'')$, for $x'' \in [x', x)$.

To finish the proof we must show that the choices of σ^x are, indeed, optimal in the interval $[x', x)$. We first prove that there exists a maximum time x'' such that π' is optimal from time x'' . Since π' is optimal from time x , there must exist a maximum time $x'' \leq x$, such that $v_k^{\pi'}(\hat{x}) = v_k(\hat{x})$ for all $\hat{x} > x''$ and all states k . Assume for the sake of contradiction that $v_k^{\pi'}(x'') \neq v_k(x'')$, for some state k . Since the choices of π' are the same throughout the interval $[x', x)$, there must be a player that can do better in the time immediately after x'' , and we get a contradiction. Hence, π' is optimal from time x'' .

From Lemma 3.1 we know that it suffices to show that (π'_1, π'_2) is a Nash equilibrium from time x' . Assume for the sake of contradiction that there exists a strategy π''_1 , a state k_0 , and a time $x_0 \in [x', x)$, such that $v_k^{\pi''_1, \pi''_2}(x'') < v_k^{\pi'_1, \pi'_2}(x'')$. Consider the finite play $\rho_{k_0, x_0}^{\pi''_1, \pi''_2} = (k_0, x_0) \xrightarrow{j_0, \delta_0} (k_1, x_1) \xrightarrow{j_1, \delta_1} \dots \xrightarrow{j_{t-1}, \delta_{t-1}} (k_t, x_t)$, and assume for simplicity that, if $x_t \geq x$, a configuration appears at time x . Let $\ell > 0$ be the minimum index such that $v_{k_{\ell'}}^{\pi''_1, \pi''_2}(x_{\ell'}) = v_{k_{\ell'}}^{\pi'_1, \pi'_2}(x_{\ell'})$ for all $\ell' \geq \ell$. Note that since the play is finite, meaning that k_t is the terminal state, equality holds for $\ell' = t$. Hence, ℓ is well-defined. Also, since (π'_1, π'_2) is optimal from time x , and x appears in a configuration of the play if $x_t \geq x$, we must have $x_\ell \leq x$. Furthermore, since there exists a time $x'' \in [x', x)$ from which π' is optimal, we may assume without loss of generality that (π'_1, π'_2) is optimal from time x_ℓ . If this is not the case we may start with a later time $x_0 \in [x_\ell, x)$, possibly with the other player. In particular, we have $v_k^{\pi''_1, \pi''_2}(x_\ell) = v_k^{\pi'_1, \pi'_2}(x_\ell) = v_k(x_\ell)$, for all states k .

Let x_ℓ be defined as above, and consider the previous transition $(j_{\ell-1}, \delta_{\ell-1})$. We may view this transition as two steps: first wait at $k_{\ell-1}$ for $\delta_{\ell-1}$ time, and then use action $j_{\ell-1}$. From the definition of ℓ we know that $v_{k_{\ell-1}}^{\pi''_1, \pi''_2}(x_{\ell-1}) < v_{k_{\ell-1}}^{\pi'_1, \pi'_2}(x_{\ell-1})$. Since π' is optimal from time x_ℓ the decrease must have occurred while waiting.

We thus have $v_{k_{\ell-1}}^{\pi''_1, \pi''_2}(x_\ell) = v_{k_{\ell-1}}^{\pi'_1, \pi'_2}(x_\ell)$, but $v_{k_{\ell-1}}^{\pi''_1, \pi''_2}(x_{\ell-1}) < v_{k_{\ell-1}}^{\pi'_1, \pi'_2}(x_{\ell-1})$. Observe that $\pi''_1(k_{\ell-1}, x'') = \lambda$ for all $x'' \in [x_{\ell-1}, x_\ell)$. It follows that:

$$\begin{aligned} v_{k_{\ell-1}}^{\pi''_1, \pi''_2}(x_{\ell-1}) &= (x_\ell - x_{\ell-1})r_{k_{\ell-1}} + v_{k_{\ell-1}}^{\pi''_1, \pi''_2}(x_\ell) = (x_\ell - x_{\ell-1})r_{k_{\ell-1}} + v_{k_{\ell-1}}^{\pi'_1, \pi'_2}(x_\ell) < v_{k_{\ell-1}}^{\pi'_1, \pi'_2}(x_{\ell-1}) \\ &\Rightarrow \epsilon r_{k_{\ell-1}} + v_{k_{\ell-1}}^{\pi'_1, \pi'_2}(x_\ell) < v_{k_{\ell-1}}^{\pi'_1, \pi'_2}(x_{\ell-1}) - \epsilon. \end{aligned} \quad (2)$$

Hence, $\lambda_{k_{\ell-1}}$ is an improving switch w.r.t. σ^x in G^{x_ℓ} . On the other hand, since σ^x is optimal for $G^{x, x-x_\ell}$ it is also optimal for $G^{x_\ell, 0}$. Hence, $x_\ell = \text{NextEventPoint}(G^x)$, and we get a contradiction from the fact that $x^0 < x_\ell$.

The case for Player 2 is identical.

Let us note that the implication in (2) does not easily work for the more general strategies described in Remark 3.3. \square

Lemma 3.7 allows us to compute optimal strategies by backward induction once the values $v_k(1)$ at time 1 are known for all states $k \in S$. Finding $v_k(1)$ and corresponding optimal strategies from time 1 is, fortunately, not difficult. Indeed, when $x = 1$ time does not increase further, and we simply solve the priced game G' that defines G . The resulting algorithm is shown in Figure 5. Note that the choice of first using the `ExtendedDijkstra` algorithm and then the `StrategyIteration` algorithm is to facilitate the analysis in Section 3.2. In fact, any algorithm for solving priced games could be used. By observing that `SolveSPTG` repeatedly applies Lemma 3.7 to construct optimal strategies by backward induction we get the following theorem.

Function SolveSPTG(G)

```

(v(1), (π1(1), π2(1))) ← ExtendedDijkstra( $G'$ );
x ← 1;
while x > 0 do
  (ax(k) + εbx(k), (σ1, σ2)) ← StrategyIteration( $G^x$ , (π1(x), π2(x)));
  x' ← NextEventPoint( $G^x$ );
  forall k ∈ S and x'' ∈ [x', x] do
    vk(x'') ← vk(x) + bx(k)(x - x'');
    π1(k, x'') ← σ1(k);
    π2(k, x'') ← σ2(k);
  x ← x';
return (v, (π1, π2));

```

Figure 5: Algorithm for solving a simple priced timed game $G = (G', (r_k)_{k \in S})$.

Theorem 3.8 *If SolveSPTG terminates, it correctly computes the value function and optimal strategies for both players.*

Note that SolveSPTG resembles the sweep-line algorithm of Shamos and Hoey [15] for the line segment intersection problem. At every time x we have n ordered sets of line segments with an intersection within one set at the next event point $x' = \text{NextEventPoint}(G^x)$. When handling the event point, the order of the line segments is updated, and we move on to the next event point.

3.2 Bounding the number of event points

Let G be an SPTG. Recall that the only difference between G^x and $G^{x'}$, for $x \neq x'$, are the costs of actions λ_k , for $k \in S$, if $v_k(x) \neq v_k(x')$. The actions available from each state are therefore the same, and a strategy profile σ for G^x can, thus, also be interpreted as a strategy profile for $G^{x'}$. To bound the number of event points we assign a potential to each strategy profile σ , such that the potential strictly decreases when one of the players performs a single improving switch. Furthermore, the potential is defined independently of the values $v_k(x)$. It then follows that the number of single improving switches performed by the SolveSPTG algorithm is at most the total number of strategy profiles for G^x . We further improve this bound to show that the number of event points is at most exponential in the number of states. This improves the previous bound by Rutkowski [14].

Let n be the number of states of G , let N be the number of distinct rates, including rate 0 for the terminal state \perp . Assume that the distinct rates are ordered such that $r_1 < r_2 < \dots < r_N$. Recall that $r(k, \sigma)$ is the rate of the waiting state reached from k in σ . Let

$$\text{count}(\sigma, i, \ell, r) = |\{k \in S_i \mid \ell(k, \sigma) = \ell \wedge r(k, \sigma) = r\}|$$

be the number of states controlled by Player i at distance ℓ from \perp in σ that reach a waiting state with rate r .

For every strategy profile σ for the priced games G^x , for $x \in (0, 1]$, define the potential $P(\sigma) \in \mathbb{N}^{n \times N}$ as an integer matrix as follows.

$$P(\sigma)_{\ell, r} = \text{count}(\sigma, 2, \ell, r) - \text{count}(\sigma, 1, \ell, r)$$

I.e., rows correspond to lengths, columns correspond to rates, and entries count the number of corresponding Player 2 controlled states minus the number of corresponding Player 1 controlled states.

We define a lexicographic ordering of potential matrices where, firstly, entries corresponding to lower rates are of higher importance. Secondly, entries corresponding to shorter lengths are more important. Formally, we write $P(\sigma) \prec P(\sigma')$ if and only if there exists ℓ and r such that:

$$\begin{aligned}
P(\sigma^{(1)}) &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} & P(\sigma^{(2)}) &= \begin{bmatrix} -1 & 0 & -1 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \\
P(\sigma^{(3)}) &= \begin{bmatrix} -1 & 1 & 0 & 1 \\ -1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} & P(\sigma^{(4)}) &= \begin{bmatrix} -1 & 0 & -1 & 1 \\ -1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}
\end{aligned}$$

Figure 6: Example of potential matrices of the strategy profiles from Figure 1.

- $P(\sigma)_{\ell',r'} = P(\sigma')_{\ell',r'}$ for all $r' < r$ and $1 \leq \ell' \leq n$.
- $P(\sigma)_{\ell',r} = P(\sigma')_{\ell',r}$ for all $\ell' < \ell$.
- $P(\sigma)_{\ell,r} < P(\sigma')_{\ell,r}$.

Figure 6 shows an example of the potential matrices of the strategy profiles shown in Figure 1. We use the following notation:

- $\sigma^{(1)}$ is the strategy profile used at time $x = 1$,
- $\sigma^{(2)}$ is the strategy profile used at time $x \in [2/3, 1)$,
- $\sigma^{(3)}$ is the strategy profile used at time $x \in [1/3, 2/3)$,
- and $\sigma^{(4)}$ is the strategy profile used at time $x \in [0, 1/3)$.

$\sigma^{(1)}$ is also shown in Figure 2. Observe that $P(\sigma^{(1)})_{1,1} = 0$ because states 1 and 5 are controlled by Player 2 and 1, respectively, and both move directly to \perp , which has rate 0. Also note that the potentials do indeed decrease for the four matrices. At each event point the strategies are updated for multiple states, however.

Lemma 3.9 *Let σ be a strategy profile that is optimal for $G^{x,0}$, for some $x \in (0, 1]$. Let $j \in A^i$ be an improving switch for Player i w.r.t. σ in the priced game G^x . Then $P(\sigma[j]) \prec P(\sigma)$.*

Proof: Consider the game G^x . Recall that for every strategy profile σ' and state $k \in S$, we let $u(k, \sigma', G^x) = a(k, \sigma', G^x) + \epsilon r(k, \sigma')$, where ϵ is an infinitesimal. We also have $u(k, \sigma', G^{x,0}) = a(k, \sigma', G^x)$. Since σ is optimal for $G^{x,0}$ we must have $a(k, \sigma) = a(k, \sigma[j])$ for all $k \in S$. Indeed, otherwise j would also be an improving switch w.r.t. σ in $G^{x,0}$, implying that σ is not optimal for $G^{x,0}$.

Let k be the state from which the action j originates. It then follows that $u(k, \sigma) \neq \infty$ and $u(k, \sigma[j]) \neq \infty$. I.e., it is not possible for exactly one of the payoffs to be infinite, and if both payoffs are infinite then j would not be an improving switch.

Assume that $i = 1$. Since $j \in A_k$ is an improving switch for Player 1 we have $\nu(k, \sigma[j]) < \nu(k, \sigma)$. It is, thus, either the case that $r(k, \sigma[j]) < r(k, \sigma)$, or that $r(k, \sigma[j]) = r(k, \sigma)$ and $\ell(k, \sigma[j]) < \ell(k, \sigma)$. In both cases the most significant entry ℓ, r for which $P(\sigma)_{\ell,r} \neq P(\sigma[j])_{\ell,r}$ is $\ell = \ell(k, \sigma[j])$ and $r = r(k, \sigma[j])$. Indeed, all states with new valuations in $\sigma[j]$ move through state k and, thus, have same rates but larger lengths. Since $i = 1$ we have $P(\sigma)_{\ell,r} < P(\sigma[j])_{\ell,r}$ and, thus, $P(\sigma[j]) \prec P(\sigma)$.

The case for $i = 2$ is similar. $j \in A_k$ is an improving switch for Player 2, implying that either $r(k, \sigma[j]) > r(k, \sigma)$, or $r(k, \sigma[j]) = r(k, \sigma)$ and $\ell(k, \sigma[j]) > \ell(k, \sigma)$. The most significant entry ℓ, r for which $P(\sigma)_{\ell,r} \neq$

$P(\sigma[j])_{\ell,r}$ is then $\ell = \ell(k, \sigma)$ and $r = r(k, \sigma)$. Since $i = 2$ we again have $P(\sigma)_{\ell,r} < P(\sigma[j])_{\ell,r}$ and subsequently $P(\sigma[j]) \prec P(\sigma)$. □

Theorem 3.10 *The total number of event points for any SPTG G with n states is $L(G) \leq \min\{12^n, \prod_{k \in S} (|A_k| + 1)\}$. Furthermore, if there is only one player, $L(G) = O(n^2)$.*

Proof: Consider the variant of the `SolveSPTG` algorithm where `StrategyIteration` only performs single improving switches for both players. I.e., when solving G^x , for some $x \in (0, 1]$, Player 1 performs one improving switch, then Player 2 repeatedly performs single improving switches as long as possible, and then the process is repeated. The resulting optimal strategy profile σ^x is then used as the starting point for solving the next priced game $G^{x'}$, for $x' = \text{NextEventPoint}(G^x)$.

Once the initial strategy profile $\sigma = (\pi_1(1), \pi_2(1))$ is found, any strategy profile σ' that is subsequently produced by the `StrategyIteration` algorithm at some time x is optimal for the priced game $G^{x,0}$. I.e., σ^x is optimal for all $G^{x''}$ with $x'' \in (x', x]$, where $x' = \text{NextEventPoint}(G^x)$. In particular, the payoffs resulting from σ^x and $\sigma^{x'}$ in $G^{x'}$ only differ by some second order term. Hence, we can apply Lemma 3.9 to the strategy profiles, and conclude that the potential decreases with every improving switch. From this we immediately get that the total number of strategy profiles in G^x , $\prod_{k \in S} (|A_k| + 1)$, is an upper bound on $L(G)$.

We next show that $L(G) \leq 12^n$. A matrix $P \in \mathbb{N}^{n \times N}$ corresponding to a legal potential can always be constructed in the following way. Let each entry (ℓ, r) be associated with a set $S_{\ell,r}$ of corresponding states. I.e., $S_{\ell,r}$ contains the states for which it takes ℓ moves to reach \perp in the priced game, and the rate encountered is the r 'th smallest rate of the game. Pick a non-empty subset of the columns $C \subseteq \{1, \dots, N\}$. This will be the columns, such that in column r , there is an ℓ such that $\text{count}(\sigma, 2, \ell, r) \neq 0$ or $\text{count}(\sigma, 1, \ell, r) \neq 0$. This can be done in at most $2^N - 1 \leq 2^{n+1}$ ways. Next, assign states to the sets of the entries. If $S_{\ell,r} \neq \emptyset$, then we must also have $S_{\ell',r} \neq \emptyset$ for all $\ell' < \ell$, by definition. This allows us to assign states to sets in an ordered way. Let (ℓ, r) be the current entry starting from $\ell = 1$ and $r = \min C$. The current entry will be lexicographic increasing in (r, ℓ) . Repeatedly add a state from either S_1 or S_2 to $S_{\ell,r}$ and update the current entry in one of the following three ways:

- Do nothing: More states will be assigned to $S_{\ell,r}$.
- Move to the next row: No more states will be assigned to $S_{\ell,r}$, but some will be assigned to $S_{\ell,r+1}$.
- Move to the beginning of the next column of C : No more states will be assigned to $S_{\ell,r'}$ for any r' .

There are n (one for each state in the game) such iterations, and in each iteration there are at most six possible options. Hence, the states can be added in at most 6^n ways. Furthermore, we do not need to update the current entry after the last state has been added, which saves us a factor of 3. The total number of possible matrices P is, thus, at most 12^n .

When there is only one player i the argument becomes much simpler. Observe that the rates change monotonically when going back in time: if $i = 1$ the rates decrease, and if $i = 2$ the rates increase. Furthermore, at every event point at least one state changes rate. Hence, there can be at most $nN \leq n(n+1)$ event points. □

Theorem 3.11 *`SolveSPTG` solves any SPTG G in time $O(m \cdot \min\{12^n, \prod_{k \in S} (|A_k| + 1)\})$ in the unit cost model, where n is the number of states and m is the number of actions. Alternatively, the variant of `SolveSPTG` that uses the `ExtendedDijkstra` algorithm instead of `StrategyIteration` solves G in time $O(L(G)(m + n \log n))$.*

Proof: The correctness of `SolveSPTG` follows from Theorems 3.8 and 3.10.

For the first bound we get from the proof of Theorem 3.10 that, in fact, not only the number of event points, but also the number of single improving switches is bounded by $\min\{12^n, \prod_{k \in S} |A_k|\}$. Valuations for a strategy profile σ can be computed in time $O(n)$, and then the next event point can be computed in time $O(m)$. I.e., for each $k \in S$ we find the next event point at time x among the intersections of $f_{\sigma(k),x}$ and $f_{j,x}$, for $j \in A_k$.

For the second bound we are using the `ExtendedDijkstra` algorithm of Khachiyan *et al.* [13] instead of `StrategyIteration` in the inner while-loop. The `ExtendedDijkstra` algorithm has the same complexity as Dijkstra's algorithm². Fredman and Tarjan [10] showed that, using Fibonacci heaps, Dijkstra's algorithm can solve the single source shortest path problem for a graph with n vertices and m edges in time $O(m + n \log n)$. \square

Theorem 3.2 follows as a corollary of Theorem 3.11, since `SolveSPTG` is always guaranteed to compute optimal strategies, and the resulting value functions are continuous piecewise linear functions.

4 Priced timed games

One-clock priced timed games (PTGs) extend SPTGs in two ways. First, actions are associated with time intervals during which they are available, and second, certain actions will cause the time to be reset to zero. Also, we do not require the time to run from zero to one.

Formally, a PTG G can be described by a tuple $G = (S_1, S_2, (A_k)_{k \in S}, (c_j)_{j \in A}, d, (r_k)_{k \in S}, (I_j)_{j \in A}, R)$, where $S = S_1 \cup S_2$ and $A = \bigcup_{k \in S} A_k$. The complete description of the individual components of G is as follows. Note that only the last two components are new compared to priced games and SPTGs.

- S_i , for $i \in \{1, 2\}$, is a set of states controlled by Player i .
- A_k , for $k \in S$, is a set of actions available from state k .
- $c_j \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, for $j \in A$, is the cost of action $j \in A$.
- $d : A \rightarrow S \cup \{\perp\}$ is a mapping from actions to destinations with \perp being the terminal state.
- $r_k \in \mathbb{R}_{\geq 0}$, for $k \in S$, is the rate for waiting at state k .
- I_j , for $j \in A$, is the *existence interval* (a real interval) of action j during which it is available.
- $R \subseteq A$ is the set of *reset actions*.

To simplify the statements of many of the remaining lemmas we let (n, m, r, d) -PTG be the class of all PTGs consisting of n states, r of which are the destination of some reset action, m actions and d distinct endpoints of existence intervals.

We let $e(I)$ be the set of endpoints of interval I , and define $M = \max_{j \in A} e(I_j)$. I.e., after time M no actions are available and the game must end. Note that PTGs are often defined with existence intervals for both states and actions. For convenience, we decided to omit this feature since it is not difficult to translate between the two version.

PTGs are played like SPTGs with the exception that using a reset transition resets the time to zero and that the actions must be available when used. We, thus, again operate with configurations $(k, x) \in S \times [0, M]$ corresponding to a pebble being placed on state k at time x . The player controlling state k chooses an action $j \in A_k$ and a delay $\delta \geq 0$, such that j is available at time $x + \delta$. I.e., $x + \delta \in I_j$. We assume for simplicity that such an action is always available. The pebble is then moved to state $d(j)$, the time is incremented to $x + \delta$ if $j \notin R$ and reset to zero otherwise, and the play continues. The game ends when the terminal state \perp is reached.

²To get this bound for the Extended Dijkstra's algorithm, actions of the maximizer should not be inserted into the priority queue. Instead, a choice of action for the maximizer for a state is fixed when the values of all possible successors of that state are known.

We again let a play be a sequence of legal steps starting from some configuration (k_0, x_0) :

$$\rho = (k_0, x_0) \xrightarrow{j_0, \delta_0} (k_1, x_1) \xrightarrow{j_1, \delta_1} \dots$$

where, for all $\ell \geq 0$, $x_\ell + \delta_\ell \in I_{j_\ell}$, and if $j_\ell \in R$ then $x_{\ell+1} = 0$. The costs of infinite plays and finite plays ending at the terminal state \perp are defined analogously to SPTGs.

Let $\text{Plays}(i)$ be the set of finite plays ending at a state controlled by Player i . Note that $\rho \in \text{Plays}(i)$ specifies the current state and time, as well as the history leading to this configuration. A (positional) strategy for Player i is again defined as a map $\pi_i : S_i \times [0, M] \rightarrow A \cup \{\lambda\}$ from configurations of the game to choices of actions. Again, for every $k \in S_i$ and $x \in [0, 1)$, if $\pi_i(k, x) = \lambda$ then we require that there exists a $\delta > 0$ such that for all $0 \leq \epsilon < \delta$, $\pi_i(k, x + \epsilon) = \lambda$. Let $\delta_{\pi_i}(k, x) = \inf\{x' - x \mid x \leq x' \leq 1, \pi_i(k, x') \neq \lambda\}$ be the delay before the pebble is moved when starting in state k at time x for some strategy π_i . Previous works have defined such strategies in other ways, see Remark 3.3.

A *history-dependent strategy* for Player i is a map $\tau_i : \text{Plays}(i) \rightarrow (A, \mathbb{R}_{\geq 0})$ that maps every play ρ ending in a state $k \in S_i$ to an action $j \in A_k$ and a delay t . We will only use history-dependent strategies in the proof of one lemma (Lemma 4.2). Note that history-dependent strategies generalize positional strategies. We denote the set of history-dependent strategies for Player i by $T_i(G)$, where G is omitted if it is clear from the context. Similarly, the set of positional strategies for Player i is denoted by $\Pi_i(G)$.

Let $\rho_{k,x}^{\tau_1, \tau_2}$ be the play generated when, starting from (k, x) , the players play according to τ_1 and τ_2 . The corresponding value function is again defined as:

$$v_k^{\tau_1, \tau_2}(x) = \text{cost}(\rho_{k,x}^{\tau_1, \tau_2}).$$

Best response, lower and upper value functions are again defined as:

$$\begin{aligned} v_k^{\tau_1}(x) &= \sup_{\tau_2 \in T_2} v_k^{\tau_1, \tau_2}(x) \\ v_k^{\tau_2}(x) &= \inf_{\tau_1 \in T_1} v_k^{\tau_1, \tau_2}(x) \\ \underline{v}_k(x) &= \sup_{\tau_2 \in T_2} v_k^{\tau_2}(x) = \sup_{\tau_2 \in T_2} \inf_{\tau_1 \in T_1} v_k^{\tau_1, \tau_2}(x) \\ \bar{v}_k(x) &= \inf_{\tau_1 \in T_1} v_k^{\tau_1}(x) = \inf_{\tau_1 \in T_1} \sup_{\tau_2 \in T_2} v_k^{\tau_1, \tau_2}(x) \end{aligned}$$

Bouyer *et al.* [8] proved the following fundamental theorem.

Theorem 4.1 (Bouyer *et al.* [8]) *For every PTG G , there exist value functions $v_k(x) := \underline{v}_k(x) = \bar{v}_k(x)$. Moreover, a player can get arbitrarily close to the values even when restricted to playing positional strategies:*

$$\begin{aligned} \underline{v}_k(x) &= \sup_{\pi_2 \in \Pi_2} \inf_{\pi_1 \in \Pi_1} v_k^{\pi_1, \pi_2}(x) = \sup_{\pi_2 \in \Pi_2} \inf_{\pi_1 \in \Pi_1} v_k^{\pi_1, \pi_2}(x) \\ \bar{v}_k(x) &= \inf_{\pi_1 \in \Pi_1} \sup_{\pi_2 \in \Pi_2} v_k^{\pi_1, \pi_2}(x) = \inf_{\pi_1 \in \Pi_1} \sup_{\pi_2 \in \Pi_2} v_k^{\pi_1, \pi_2}(x) \end{aligned}$$

For the purpose of solving PTGs it, thus, suffices to consider positional strategies. In the remainder of this section we will therefore restrict ourselves to positional strategies unless otherwise specified.

A strategy $\pi_i \in \Pi_i$ is ϵ -optimal for Player i for $\epsilon \geq 0$ if:

$$\forall k \in S, x \in [0, M] : |v_k^{\pi_i}(x) - v_k(x)| \leq \epsilon.$$

Since PTGs have value functions, ϵ -optimal strategies always exist for both players, for any $\epsilon > 0$. Optimal strategies do not always exist, as shown by Bouyer *et al.* [8]. Indeed, consider the PTG shown in Figure 7. State 1 is controlled by Player 1, the minimizer, and state 2 is controlled by Player 2, the maximizer. The value functions are shown on the right. Two actions leading to the terminal state are available from state 2 at time 0 and 1, respectively. Since the rate of state 2 is 0, Player 2 picks the more expensive action with

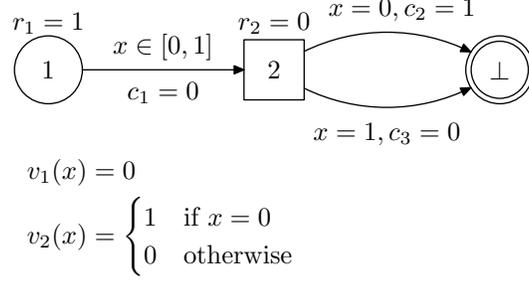


Figure 7: Example of a PTG with no optimal strategy profile.

cost $c_2 = 1$ at time 0, and at other times Player 2 waits until time 1 and picks the cheaper action with cost $c_3 = 0$. From state 1 exactly one action is available at all times, and since the rate is 1, Player 1 leaves the state *as soon as possible*, only not at time 0. Since no strategy can implement leaving as soon as possible there is no optimal strategy for Player 1. More precisely, for every waiting time δ chosen by Player 1 at time 0, there exists a smaller waiting time $\delta' < \delta$ that achieves a better value.

We reduce solving any PTG to solving a number of SPTGs. The first step towards this goal is to remove reset actions by extending the game.

Lemma 4.2 *Let G be a (n, m, r, d) -PTG. Solving G can be reduced to solving $r + 1$ $(n, m, 0, d)$ -PTGs.*

Proof: Let $\pi = (\pi_1, \pi_2)$ be any strategy profile, and suppose the play ρ_{k_0, x_0}^π is using two reset actions $j, j' \in R$ leading to the same state $d(j) = d(j') = k$. Then the configuration $(k, 0)$ appears twice in ρ_{k_0, x_0}^π , and since strategies are history-independent it appears an infinite number of times. It follows that $v_{k_0}^\pi(x_0) = \infty$. By the pigeonhole principle we get that if a play ρ_{k_0, x_0}^π uses $r + 1$ reset actions, then some state is visited twice by some reset actions, and therefore $v_{k_0}^\pi(x_0) = \infty$.

Thus, when playing G we may augment configurations by the number of times a reset action has been used, and once this number reaches $r + 1$ we may assume without loss of generality that the value is infinite. This defines a new PTG G' with states $S' = S \times \{0, \dots, r\}$ and actions $A' = A \times \{0, \dots, r\}$ in the following natural way. For $j \in A$ and $\ell \in \{0, \dots, r\}$, destinations and costs are defined as

$$d'(j, \ell) = \begin{cases} (d(j), \ell + 1) & \text{if } j \in R \text{ and } \ell < r \\ \perp & \text{if } j \in R \text{ and } \ell = r \\ (d(j), \ell) & \text{otherwise} \end{cases}$$

$$c'_{(j, \ell)} = \begin{cases} \infty & \text{if } j \in R \text{ and } \ell = r \\ c_j & \text{otherwise} \end{cases}$$

while rates, existence intervals and reset actions are the same as for the corresponding states and actions of G . Plays and value functions of G' will be denoted by ρ' and v' , respectively. We will show that for all $(k, x) \in S \times [0, M]$, $v'_{(k, 0)}(x) = v_k(x)$.

Every strategy profile π' for G' can be interpreted as a history-dependent strategy profile for G in the following way: For every play that can be achieved by moving according to π' make the corresponding choice in π' , for other plays make arbitrary choices. Also, every positional strategy profile π for G can be interpreted as a strategy profile for G' by using the same choices regardless of the number of encountered reset actions. With these interpretations we see that $\Pi_i(G) \subseteq \Pi_i(G') \subseteq T_i(G)$.

For all configurations $(k, x) \in S \times [0, M]$, if $\rho'_{(k, 0), x}^{\pi'}$ uses at most r reset actions, then $\text{cost}(\rho'_{(k, 0), x}^{\pi'}) = \text{cost}(\rho_{k, x}^\pi)$, since the actions encountered in the two games have the same costs. If $\rho'_{(k, 0), x}^{\pi'}$ uses more than r reset actions, then $\text{cost}(\rho'_{(k, 0), x}^{\pi'}) = \infty \geq \text{cost}(\rho_{k, x}^\pi)$. Hence, we always have $v'_{(k, 0)}(x) \geq v_k(x)$. Using

Theorem 4.1 it follows that:

$$\begin{aligned} v'_{(k,0)}(x) &= \inf_{\pi'_1 \in \Pi'_1} \sup_{\pi'_2 \in \Pi'_2} v'^{\pi'_1, \pi'_2}_{(k,0)}(x) \geq \inf_{\pi'_1 \in \Pi'_1} \sup_{\pi'_2 \in \Pi'_2} v_k^{\pi'_1, \pi'_2}(x) \\ &= \inf_{\pi_1 \in \Pi_1} \sup_{\pi_2 \in \Pi_2} v_k^{\pi_1, \pi_2}(x) = v_k(x) \end{aligned}$$

The first inequality follows from the costs being larger in G' , and the next equality follows Theorem 4.1; the same values can be obtained using only positional strategies in G .

Next we show that $v'_{(k,0)}(x) \leq v_k(x)$, implying that $v'_{(k,0)}(x) = v_k(x)$. This is clearly true if $v_k(x) = \infty$, thus, we may assume that $v_k(x) < \infty$. In particular, ϵ -optimal strategies do not generate plays with more than r reset actions in neither G nor G' . We see that:

$$\begin{aligned} v_k(x) &= \inf_{\pi_1 \in \Pi_1} \sup_{\pi_2 \in \Pi_2} v_k^{\pi_1, \pi_2}(x) = \inf_{\pi'_1 \in \Pi'_1} \sup_{\pi'_2 \in \Pi'_2} v_k^{\pi'_1, \pi'_2}(x) \\ &= \inf_{\pi'_1 \in \Pi'_1} \sup_{\pi'_2 \in \Pi'_2} v'^{\pi'_1, \pi'_2}_{(k,0)}(x) = v'_{(k,0)}(x) \end{aligned}$$

For the second equality we use Theorem 4.1; the values do not change even if certain history-dependent strategies are available. For the third equality we use the assumption that the values are finite. This implies that for relevant strategy profiles the values of the two games are the same.

We now know that in order to find the value $v_k(x)$ in G it suffices to find $v'_{(k,0)}(x)$ in G' . To do this we exploit the special structure of G' . We observe that states $(k, \ell) \in S \times \{0, \dots, r\}$ do not depend on states (k, ℓ') with $\ell' < \ell$. Thus, the game can be solved using backward induction on ℓ . In particular, when $v'_{(k, \ell+1)}(x)$ is known for all k and x , then the subgame consisting of states (k, ℓ) , for $k \in S$, can be viewed as an independent PTG with no reset actions. I.e., reset actions lead to states with known values, and can, thus, be thought of as going directly to the terminal state with an appropriate cost. Each subgame has n states and m actions, and there are $r + 1$ such subgames. □

Now we just need to show how to solve PTGs without resets using SPTGs.

We will show the statement using 3 reductions. First we will reduce PTGs without resets to the subclass of such games, where, for each action $j \in A$, we have $I_j \in \{(0, 1), [1, 1]\}$. Afterwards we will reduce further to the subclass of PTGs where for each action $j \in A$, we have $I_j \in \{[0, 1], [1, 1]\}$. At the end we will reduce those to SPTGs.

Let X be the set which consists of 0 and the endpoints of existence intervals of G . Let the i 'th largest element in X be M_i . Note that $M_1 = M$.

We will now define some functions on PTGs. For a PTG $G = (S_1, S_2, (A_k)_{k \in S}, c, d, (r_k)_{k \in S}, (I_j)_{j \in A}, R)$, where $R = \emptyset$, a $x \in \mathbb{R}$ and a vector $v \in (\mathbb{R}_{\geq 0 \cup \{\infty\}})^n$, let the priced game $G^{v,x} = (S_1, S_2, (A'_k)_{k \in S}, c', d')$ be defined by:

$$\begin{aligned} \forall k \in S : A'_k &= \{j \in A_k \mid x \in I_j\} \cup \{\perp_k\} \\ \forall j \in A'_k : c'_j &= \begin{cases} v_k & \text{if } j = \perp_k \\ c_j & \text{otherwise} \end{cases} \\ \forall j \in A'_k : d'_j &= \begin{cases} \perp & \text{if } j = \perp_k \\ d_j & \text{otherwise} \end{cases} \end{aligned}$$

The game $G^{v,x}$ is similar to the priced game defined in Definition 3.4. The intuition is that $G^{v,x}$ can model a specific moment in time.

Definition 4.3 For a given PTG $G = (S_1, S_2, (A_k)_{k \in S}, c, d, (r_k)_{k \in S}, (I_j)_{j \in A}, R)$, a $x \in \mathbb{R}$ and a vector $v \in (\mathbb{R}_{\geq 0 \cup \{\infty\}})^n$, let the SPTG $G^{v,x,d} = (S_1, S'_2, (A'_k)_{k \in (S_1 \cup S'_2)}, c', d', (r'_k)_{k \in (S_1 \cup S'_2)})$ be defined by:

Function SolvePTG(G)

```

 $v(M_1) \leftarrow \text{ExtendedDijkstra}((S_1, S_2, (\{j \in A_k \mid M_1 \in I_j\})_{k \in S}, c, d));$ 
 $i \leftarrow 1;$ 
while  $M_i > 0$  do
   $i \leftarrow i + 1;$ 
   $x \leftarrow \frac{M_{i-1} + M_i}{2};$ 
   $v' \leftarrow \text{ExtendedDijkstra}(G^{v(M_{i-1}), x});$ 
   $v^*(x) \leftarrow \text{SolveSPTG}(G^{v', x, M_{i-1} - M_i});$ 
  forall  $x \in (M_i, M_{i-1})$  do
     $v(x) \leftarrow v^*(\frac{x - M_i}{M_{i-1} - M_i});$ 
   $v(M_i) \leftarrow \text{ExtendedDijkstra}(G^{v^*(0), M_i});$ 
return  $v;$ 

```

Figure 8: Algorithm for solving PTGs without reset actions.

$$\begin{aligned}
S'_2 &= S_2 \cup \{\max\} \\
\forall k \in S : A'_k &= \{j \in A_k \mid x \in I_j\} \cup \{\perp_k\} \\
A_{\max} &= \{\perp_{\max}\} \\
\forall j \in A'_k : c'_j &= \begin{cases} 0 & \text{if } k = \max \\ v_k & \text{if } k \neq \max \text{ and } j = \perp_k \\ c_j & \text{otherwise} \end{cases} \\
\forall j \in A'_k : d'_j &= \begin{cases} \max & \text{if } k \in S_1 \text{ and either } j = \perp_k \\ & \text{or } d_j = \perp \\ \perp & \text{if } k \in S_2 \text{ and either } j = \perp_k \\ & \text{or } d_j = \perp \\ \perp & \text{if } k = \max \\ d_j & \text{otherwise} \end{cases} \\
\forall k \in S : r'_k &= r_k \cdot d \\
r'_{\max} &= \max_{k \in S} \{r_k\}
\end{aligned}$$

The game $G^{v,x,d}$ is constructed from the proof of Lemma 4.5, Lemma 4.6 and Lemma 4.7. The intuition is that the game can model an arbitrary length interval, in the original game, where no action changes status between available and unavailable.

Theorem 4.4 *The algorithm in Figure 8 correctly solves Priced Timed Games without reset actions.*

The proof of correctness is that the algorithm is a formalization of the reductions in Lemma 4.5, Lemma 4.6 and Lemma 4.7. Note that instead of $\frac{M_i + M_{i-1}}{2}$, any arbitrary point inside (M_i, M_{i-1}) would work.

Let $PTG_I^{n,m}$ be the subclass of PTGs, consisting of n states, m actions, none of which are reset actions, and where the existence interval for each action, j , is either $I_j = I$ or $I_j = [1, 1]$. In the latter case $d_j = \perp$. Note that for such games we can WLOG assume that $m \leq 2n^2$, because for all actions with the same existence interval, only the one with the best cost will be used.

Lemma 4.5 Any game G in $(n, m, 0, d)$ -PTG can be solved in time $O((n \log n + \min(m, n^2))d)$ using at most d calls to an oracle, R , that solves $PTG_{(0,1)}^{n,m+n}$.

We sketch the proof. It is easy to find the value of $k \in S$ at time M_1 in a priced timed game without reset actions, because no player can wait and hence the game is equivalent to a priced game. Between time M_2 and M_1 the game is nearly an SPTG, since we can simply translate by decreasing all times with M_2 and divide the times by $M_1 - M_2$ to get a game between 0 and 1 instead. After finding the value between M_2 and M_1 we can then find the value at M_2 , since we know the cost if we wait (it becomes $\lim_{x \rightarrow M_2^+} v(k, x)$), by viewing the game as a priced game at that point. We can then find the value between M_3 and M_2 , then at time M_3 and so on, until we have solved the game.

Proof: We can find $v(k, M_1)$ as the value of state k in the priced game which consists of the same states as G and the actions available at time M_1 . We can do so, because the game contains no reset actions and we can therefore neither increase nor decrease time. Note that if multiple actions, j , in A_k and $d_j = \ell$ exists for $k, \ell \in S$, we can ignore all but the one with the best cost for the controller of k . Hence we can solve such a priced game in time $O(n \log n + \min(m, n^2))$.

We now want to find $\forall k \in S, x \in (M_2, M_1) : v(k, x)$. We see that if we wait until M_1 in some state, k , the rest of the path to \perp costs $v(k, M_1)$, if we play optimally from M_1 . We see that if we start at a time x , we can not reach a time before x , because there are no reset actions. Hence, look at a modified game G' , with value function v' : G' consists of the same set of states as G , but it only has the actions available in the interval (M_2, M_1) , which, in G' , only exists in that interval, and for each state, k , an action to \perp of cost $v(k, M_1)$ which is only available at time M_1 . We will also modify G' such that we subtract M_2 from all points in time. Clearly that will not matter for plays starting after time M_2 . Note that all intervals for actions are either $(0, M_1 - M_2)$ or $[M_1 - M_2, M_1 - M_2]$. We can also divide all points in time with $M_1 - M_2$, by also multiplying the rate of each state with $M_1 - M_2$. Hence all existence intervals either have the form $(0, 1)$ or $[1, 1]$ and we clearly have that

$$\forall x \in (M_2, M_1), k \in S : v(k, x) = v'(k, \frac{x - M_2}{M_1 - M_2}).$$

We can solve G' using a call to R .

We will now find $v(k, M_2)$. If it is optimal to wait at time M_2 in state k , we have that $v(k, M_2) = \lim_{x \rightarrow 0^+} v'(k, x) = v'(k, 0)$, because we might as well wait as little as possible and then play optimally from there. Hence, $v(k, M_2)$ is the value of state k in the priced game G'' , consisting of the same states as G and the same actions as those available at time M_2 in G and for each state k , a action from k to \perp of cost $v'(k, 0)$. Like we did for M_1 we can ignore all but one action from a state to another. Hence we can solve such a priced game in time $O(n \log n + \min(m, n^2))$.

We now want to find $\forall k \in S, x \in (M_3, M_2) : v(k, x)$. We can therefore do like we did for $\forall k \in S, x \in (M_2, M_1) : v(k, x)$. Also to find $v(k, M_3)$ we can do like we did for $v(k, M_2)$. We keep on doing this until we are done.

Hence, we use d calls to R and solve $d + 1$ priced games. □

We will now to reduce a game in $PTG_{(0,1)}^{n,m}$, with value function, v , to a game in $PTG_{[0,1]}^{n,m}$ using $O(n \log n + \min(m, n^2))$ time. First note that it is easy to find $v(k, 1)$ using a priced game, because time can not change at time 1. It is clear that $v(k, 0) = \lim_{x \rightarrow 0^+} v(k, x)$, because the only option at time 0 is to wait. Hence we only need to look at finding $v(k, x)$ for $x \in (0, 1)$. To that we will use the following lemma. Note that the game G' mentioned in the lemma is in $PTG_{[0,1]}^{n,m}$.

Lemma 4.6 Let G be a $PTG_{(0,1)}^{n,m}$, with value function v . Let G' be the modified version of G , where all existence intervals of the form $(0, 1)$ in G instead have the form $[0, 1]$. Let v' be the value function for G' . We then have: $\forall k \in S, x \in (0, 1) : v(k, x) = v'(k, x)$.

Proof: Let $\epsilon > 0$. We will show that $\forall k \in S, x \in (0, 1) : v(k, x) = v'(k, x)$ by constructing a strategy, σ_1 , for player 1 that guarantees at most $v'(k, x) + \epsilon$, in G , for any $k \in S$ and for any $x \in (0, 1)$. Similarly we

will construct a strategy, σ_2 , for player 2 that guarantees at least $v'(k, x) - \epsilon$, in G , for any $k \in S$ and for any $x \in [0, 1)$.

Let σ_1' be a $\epsilon/2$ -optimal strategy for player 1 in G' . Let $r_{\max} = \max_{k \in S} r(s)$. Let σ_1^G be the optimal strategy in the priced game which consists of the same states as G , but only those actions available at time 1. Let $\sigma_1^{G'}$ be the optimal strategy in the priced game which consists of the same states as G' . It is clear that if the existence interval of $\sigma_1^{G'}(k)$ in G is $[1, 1]$ then $\sigma_1^{G'}(k) = \sigma_1^G(k)$.

We will now construct σ_1 .

$$\sigma_1(k, x) = \begin{cases} \lambda & \text{if } x = 0 \\ \sigma_1'(k, x) & \text{if } 0 < x < 1 - \frac{\epsilon}{2r_{\max}} \\ \sigma_1^{G'}(k) & \text{if } 1 - \frac{\epsilon}{2r_{\max}} \leq x < 1 \text{ and the existence} \\ & \text{interval for } \sigma_1^{G'}(k) \text{ in } G \text{ is } (0, 1) \\ \lambda & \text{if } 1 - \frac{\epsilon}{2r_{\max}} \leq x < 1 \text{ and the existence} \\ & \text{interval for } \sigma_1^{G'}(k) \text{ in } G \text{ is } [1, 1] \\ \sigma_1^G(k) & \text{if } x = 1 \end{cases}$$

Let $k \in S, x \in (0, 1)$. We will first show that $v'(k, x) \geq v'(k, 1)$. If $k \in S_2$ player 2 could simply wait until time 1, and since $r(k) \geq 0$ the statement follows. If $k \in S_1$ player 1 must keep the play away from S_2 (because that would reduce it to the first case) and have no advantages in waiting, since no new actions become available. But since all actions are available at time 1, player 1 could follow the same strategy, as he uses at time x , and get the same cost.

We will now show that σ_1 guarantees at most $v'(k, x) + \epsilon$, for $k \in S, x \in (0, 1)$ in G . We will do so by contradiction. Assume not. Hence there is a strategy σ_2 , a $x \in (0, 1)$ and a $k \in S$, such that $v_k^{(\sigma_1, \sigma_2)}(x) > v'(k, x) + \epsilon$. Let ρ be the play defined by (σ_1, σ_2) .

There are now two cases. Either $x \geq 1 - \frac{\epsilon}{2r_{\max}}$ or not.

If $x \geq 1 - \frac{\epsilon}{2r_{\max}}$, we know that

$$v_k^{(\sigma_1, \sigma_2)}(x) = \text{cost}(\rho) = \sum_{i=0}^{t-1} (\delta_i r_{k_i} + c_{j_i})$$

We have that $\sum_{i=0}^{t-1} \delta_i$ is at most $1 - x$, because there are no reset actions, $r_{k_i} \leq r_{\max}$, by definition, and $\sum_{i=0}^{t-1} c_{j_i} \leq v'(k, 1)$, by construction of σ_1 .

Hence

$$\begin{aligned} v_k^{(\sigma_1, \sigma_2)}(x) &\leq (1 - (1 - \frac{\epsilon}{2r_{\max}}))r_{\max} + v'(k, 1) \\ &= \epsilon/2 + v'(k, 1) \leq \epsilon/2 + v'(k, x) \end{aligned}$$

That is a contradiction.

Otherwise, if $x < 1 - \frac{\epsilon}{2r_{\max}}$, there are two cases. Either the play defined by (σ_1, σ_2) at some point waits until time $x' \geq 1 - \frac{\epsilon}{2r_{\max}}$ or not. If not, then the play cost at least $v'(k, x) + \epsilon/2$ because player 1 has at all times followed a strategy that guarantees at least that.

Otherwise, we can divide ρ up in two. ρ^1 is the first part. The second part, ρ^2 begins in some state k' and at time x' such that $x' = 1 - \frac{\epsilon}{2r_{\max}}$. Note that this might be in the middle of a wait period. Clearly $\text{cost}(\rho) = \text{cost}(\rho^1) + \text{cost}(\rho^2)$. We must have that $\text{cost}(\rho^1) + v'(k', x') \leq v'(k, x) + \epsilon/2$, because we followed a ϵ optimal strategy for player 1 in G' in ρ^1 . By the first part we also know that $\text{cost}(\rho^2) \leq v'(k', x') + \epsilon/2$.

Hence it is easy to see that $\text{cost}(\rho) \leq v'(k, x) + \epsilon$. That is a contraction.

The construction of σ_2 can be done symmetrically. □

Lemma 4.7 Solving any game G in $G' \in PTG_{[0,1]}^{n,m}$ can be polynomially reduced to solving an SPTG with $n + 1$ states and $m + 1$ actions.

Proof: Player 2 will never use a action, j , to \perp except at time 1 in a simple priced timed game, because player 2 might as well wait until time 1 before using j , which will not decrease the cost because rates are non-negative. Hence we can change all actions, j , of the form $[1, 1]$ to $[0, 1]$ if $j \in A_k, k \in S_2$, without changing the value functions. We will create a new state, max, with the maximum rate in the game, belonging to player 2, which has a action to \perp of cost 0 and existence interval $[0, 1]$. We will now redirect all actions which have existence interval $[1, 1]$ to max and change the existence interval to $[0, 1]$. We can see that player 1 will only use the actions to max at time 1, since it is cheaper to wait to time 1 and then move to max.

Now all existence intervals have the form $[0, 1]$. It is easy to see that we only need one action, j , for $j \in A_k$ and $d_j = \ell$ for any pair $k, \ell \in S$, because the controller of k will, when playing optimally, only use the action with the best, for that player, cost. Hence the game is a simple priced timed game. \square

Lemma 4.8 Any game G in (n, m, r, d) -PTG, can be solved in time $O((r + 1)d(n \log n + \min(m, n^2)))$ using at most $(r + 1)d$ calls to an oracle R that solves SPTGs with $n + 1$ states and at most $m + n + 1$ actions.

Proof: The proof is a simple consequence of Lemma 4.2, Lemma 4.5, Lemma 4.6 and Lemma 4.7. \square

Note that d is bounded by $2m + 1$ and r is bounded by n .

Theorem 4.9 Any game G in (n, m, r, d) -PTG, can be solved in time

$$O((r + 1)d(\min(m, n^2) + n \cdot \min\{12^n, \prod_{k \in S} (A_k + 1)\})).$$

Proof: The proof is a consequence of Theorem 3.11 and Lemma 4.8. Note that we only get $\prod_{k \in S} (A_k + 1)$ and not $\prod_{k \in S} (A_k + 2)$, because the additional actions we add to each state (using Definition 4.3 and 3.4) both goes to \perp and hence we only need one of them. \square

Theorem 4.10 The BCFL-ABM algorithm solves any PTG G using at most

$$m \cdot n^{O(1)} \min\{12^n, \prod_{k \in S} (A_k + 1)\}$$

iterations.

Proof: Note that Lemma 4.9 gives us an upper bound on the number of line segments of the value functions of G , because the number of line segments is a lower bound of the size of the output. By Bouyer *et al.* [8], page 11, we know that the number of iterations needed for the BCFL-ABM algorithm is at most the number of line segments times n . \square

Theorem 4.11 Any priced timed game, G , in (n, m, r, d) -PTG, where all states have rate 1 and all actions have cost 0, can be solved in time $O((r + 1)d(n \log n + \min(m, n^2)))$.

Proof: If we use Lemma 4.2, Lemma 4.5, Lemma 4.6 and Lemma 4.7 on such a game, we get $(r + 1)d$ SPTGs. If we look carefully at the lemmas we see that all states, in the SPTGs have rate c , for some $c > 0$. c depends on the interval length. Also, all actions that do not go to \perp or max have cost 0.

We now need to bound the number of event points. We will show that $L(G') = 1$ for G' being any of the SPTGs generated.

Look at the priced game G^1 , as defined in section 3.1. Let σ be some optimal strategy profile for G^1 . We see that, if $r(k, \sigma) = 0$ for some $k \in S$, we can not have passed through any states in S_2 , since it is optimal

to wait until time 1 for player 2. Since all actions of positive cost either goes to a state in S_2 or from a state in S_2 , we must have that $v_k^\sigma = 0$.

For convenience we repeat the definition of the next event point and the function f here. The definition of the next event point, $\text{NextEventPoint}(G^x)$, was:

$$\max \{0\} \cup \{x' \in [0, x) \mid \exists k \in S, j \in A_k : f_{j,x}(x) \neq f_{\pi(k),x}(x) \wedge f_{j,x}(x') = f_{\pi(k),x}(x')\}.$$

The definition of f was

$$f_{j,x}(x'') = c_j + a^x(d(j)) + b^x(d(j))(x - x'').$$

Note that $b^x(d(j))$ corresponds to the rate of the next state we wait in if both players follow σ and $f_{j,x}(x'')$ is the cost to reach \perp if both players follow σ . Hence $f_{j,x} \geq 0$ and if $b^1 = 0$ then $f_{j,x}(x'') = 0$ by the preceding, because σ was optimal. Note that if $f_{j,1}(1) \neq f_{\pi(k),1}(1)$, then at least the larger expression of the two must have $b^1(d(j)) = c$ and therefore we have that for all $x \in [0, 1) : f_{j,1}(x) \neq f_{\pi(k),1}(x)$, because either the $b^1(d(j))$'s are equal in the two expressions, in which case the difference between the two expressions do not change with x , or one is positive and the other is 0 for all $x \in [0, 1]$.

Therefore we can apply Theorem 3.11 and get that SolveSPTG solves any of $(r+1)d$ SPTGs in time $O(m + n \log n)$. We therefore solve all in time $O((r+1)d(\min(m, n^2) + n \log n))$. The reductions also required time $O((r+1)(n \log n + \min(m, n^2))d)$. Hence, our time bound becomes $O((r+1)d(n \log n + \min(m, n^2)))$. \square

Theorem 4.12 *Any priced timed automata (i.e., all states are controlled by Player 1), G in (n, m, r, d) -PTG, can be solved in time $O((r+1)dn^2(\min(m, n^2) + n \log n))$.*

Proof: If we use Lemma 4.2, Lemma 4.5 and Lemma 4.6 on a priced timed automata, we get $(r+1)d$ priced timed games, without resets where all existence intervals are either $[0, 1]$ or $[1, 1]$ and all states belong to Player 1. The algorithm described in Figure 5, solves SPTGs by first solving them for time 1, as a priced game, and then solve them by induction backwards through time. We can still solve the game at time 1, and there is no differences in the induction, since no actions become available at time x for $x < 1$. Hence, from Theorem 3.10 and Theorem 3.11 we get that we can solve such games in time $n^2(\min(m, n^2) + n \log n)$. The reductions also required time $O((r+1)d(\min(m, n^2) + n \log n))$ and therefore the total time complexity is $O((r+1)dn^2(\min(m, n^2) + n \log n))$. \square

5 Concluding remarks

We have presented an algorithm for solving one clock priced timed games with a complexity which is close to linear in L , with $L = L(G)$ being a lower bound on the size of the object to be produced as output. We think it is an attractive candidate for implementation.

We have also given a new upper bound on L . While it is better than previous bounds, we do not expect this bound to be optimal. It seems to be a ‘‘folklore theorem’’ that L does not become very big for games arising in practice. We would like to suggest the following conjecture.

Conjecture 5.1 *For all SPTGs G , $L(G) \leq p(n)$ for some polynomial p .*

Note that if this conjecture is established, it implies that our algorithm as well as the BCFL-ABM algorithm runs in time polynomial in the size of its input.

6 Acknowledgements

We would like to thank Kim Guldstrand Larsen for many helpful discussions and comments.

References

- [1] R. Alur, M. Bernadsky, and P. Madhusudan. Optimal reachability for weighted timed games. In *Proc. of 31st ICALP*, pages 122–133, 2004.
- [2] R. Alur and D. L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:183–235, 1994.
- [3] R. Alur, S. L. Torre, and G. J. Pappas. Optimal paths in weighted timed automata. In *Proc. of 4th HSCC*, pages 49–62, 2001.
- [4] G. Behrmann, A. Fehnker, T. Hune, K. G. Larsen, P. Pettersson, J. M. T. Romijn, and F. W. Vaandrager. Minimum-cost reachability for priced timed automata. In *Proc. of 4th HSCC*, pages 147–161, 2001.
- [5] P. Bouyer, T. Brihaye, V. Bruyere, and J.-F. Raskin. On the optimal reachability problem. *Formal Methods in System Design*, 31(2):135 – 175, 2007.
- [6] P. Bouyer, T. Brihaye, and N. Markey. Improved undecidability results on weighted timed automata. *Inf. Process. Lett.*, 98:188–194, 2006.
- [7] P. Bouyer, F. Cassez, E. Fleury, and K. G. Larsen. Optimal strategies in priced timed game automata. In *Proc. of 24th FSTTCS*, pages 148–160, 2004.
- [8] P. Bouyer, K. G. Larsen, N. Markey, and J. I. Rasmussen. Almost optimal strategies in one clock priced timed games. In *Proc. of 26th FSTTCS*, pages 345–356, 2006.
- [9] T. Brihaye, V. Bruyere, and J.-F. Raskin. On optimal timed strategies. In *Proc. of 3rd FORMATS*, pages 49–64, 2005.
- [10] M. L. Fredman and R. E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34(3):596–615, 1987.
- [11] R. Howard. *Dynamic programming and Markov processes*. MIT Press, 1960.
- [12] M. Jurdzinski and A. Trivedi. Reachability-time games on timed automata. In *Proc. of 34th ICALP*, pages 838–849, 2007.
- [13] L. Khachiyan, E. Boros, K. Borys, K. Elbassioni, V. Gurvich, G. Rudolf, and J. Zhao. On short paths interdiction problems: Total and node-wise limited interdiction. *Theory of Computing Systems*, 43:204–233, 2008.
- [14] M. Rutkowski. Two-player reachability-price games on single-clock timed automata. In *Proc. of 9th QAPL*, pages 31–46, 2011.
- [15] M. I. Shamos and D. Hoey. Geometric intersection problems. In *Proc. of 17th FOCS*, pages 208–215, 1976.
- [16] S. L. Torre, S. Mukhopadhyay, and A. Murano. Optimal-reachability and control for acyclic weighted timed automata. In *Proc. of 2nd TCS*, pages 485–497, 2002.