Average-Price-per-Reward Games on Hybrid Automata with Strong Resets*

Marcin Jurdziński, Ranko Lazić, and Michał Rutkowski

Department of Computer Science, University of Warwick, UK

Abstract. We study price-per-reward games on hybrid automata with strong resets. They generalise priced games previously studied and have applications in scheduling. We obtain decidability results by a translation to a novel class of finite graphs with price and reward information, and games assigned to edges. The cost and reward of following an edge are determined by the outcome of the edge game that is assigned to it.

1 Introduction

Hybrid systems and automata. Systems that exhibit both discrete and continuous behaviour are referred to as *hybrid systems*. Continuous changes to the system's state are interleaved with discrete ones, which may alter the constraints for future continuous behaviours. *Hybrid automata* are a formalism for modeling hybrid systems [1]. Hybrid automata are finite automata augmented with continuous real-valued variables. The discrete states can be seen as modes of execution, and the continuous changes of the variables as the evolution of the system's state over time. The mode specifies the continuous dynamics of the system, and mode changes are triggered by the changes in variable's values.

Reachability [2–5] and optimal reachability [6, 7] analysis for hybrid automata have been studied. In [6, 8] the optimality of infinite behaviours is also addressed.

Optimal schedule synthesis. Hybrid systems have been successfully applied to modeling scheduling problems [9]. In this setting, an execution of the automaton is a potential schedule. In [8], the authors equip timed automata, a subclass of hybrid systems, with price and reward information. Each schedule comes at a price, but provides a certain reward. The price-over-reward ratio can be seen as a measure of how cost-effective the schedule is. A natural example of a reward is time. In this case, price-per-time unit is being optimised. The problem that arises is the synthesis of an optimal schedule, i.e., a schedule that minimises the price-over-reward ratio. Reachability-price-per-reward analysis is used in the synthesis of finite optimal schedules. When dealing with reactive behaviour, optimality of infinite schedules becomes more important. Average-price-per-reward analysis, where the average price-over-reward ratio of a single step in the execution is optimised, is used in this context [8].

* This research was supported in part by EPSRC project EP/E022030/1.

We follow this direction and study the problem in the context of hybrid automata with strong resets. Our research shares the same motivation, but both the model and the techniques used differ. In [8] timed automata, a different class of hybrid automata, is considered, and an abstraction technique, known as "corner-point abstraction", is used. We, on the other hand, use an abstraction, that was first proposed in [6], to reduce to price-reward graphs, that are introduced in this paper.

Controller synthesis. The designer of a system often lacks full control over its operation. The behaviour of the system is a result of an interaction between a controller and the environment. This gives rise to the *controller synthesis* problem (first posed by Church [10]), where the goal is to design a program such that, regardless of the the environment's behaviour, the system behaves correctly and optimally. A game-based approach was proposed in [11], and was applied to hybrid automata [12, 13] and timed automata [14]. There are two players, *controller* and *environment*, and they are playing a zero-sum game. The game is played on the hybrid automaton and consists of rounds. As usual, we use player Min to denote the controller and player Max to denote the environment. In each round, Min proposes a transition. In accordance with the game protocol, Max can choose to perform this or another transition.

Determinacy and decidability are important properties of zero-sum games. A determined zero-sum game has a value, and admits almost-optimal controllers (strategies). A determined game is decidable if, given some some rational number, we can decide whether the value of the game is greater than the number.

Hybrid games with strong resets. We are considering a subclass of hybrid automata: hybrid automata with strong resets (HASR). In order to represent the automaton finitely, we require that all the components of the system are first-order definable over the ordered field of reals. The term "strong resets" comes from the property of the system that all continuous variables are nondeterministically reset after each discrete transition. As opposed to timed automata, where flow rates are constant, and reseting of the variables upon a discrete transition is not compulsory [2], HASR allow for rich continuous dynamics [4, 13, 12].

In the game setting, we allow only for alternating sequences of timed and discrete transitions [12, 13]. A timed transition followed by a discrete one will be called a timed action. Allowing an arbitrary number of continuous transitions prior to a discrete one, without the requirement of o-minimality, renders it impossible to construct a bisimulation of finite index [15, 16].

Contributions. We are considering *average-price-per-reward games*, where players are trying to optimise the average price-over-reward ratio of a timed action. Our main result is that average-price-per-reward games are determined and decidable. It is obtained through a reduction to games on finite price-reward graphs (PRGs) introduced in this paper.

To reduce hybrid average-price-per-reward games to their counterparts on PRGs we use the same equivalence as in [6]. However, there are two significant contributions with respect to [6]. Firstly, we are considering the average price-over-reward ratio, whereas only average price per transition was considered in [6]. The first is significantly more complex. Secondly, we introduce a novel class of finite graphs with price and reward information, and games assigned to edges (PRG). In this paper we show that average-price-per-reward games on PRGs are determined and decidable.

We believe that our results and technical developments concerning PRGs are interesting in their own right. To characterise game values we use a technique, referred to as optimality equations [14, 6]. What is novel is that we use the values of edge games to express optimality criteria in these equations. The proof that solutions to the optimality-equations exist (and hence the games are determined) relies on the properties of the equations, not of a particular game (on a PRG). This makes us believe that our technique is robust, and can be used to solve related games such as, reachability-price-per-reward. To show determinacy and decidability we only need to express optimality criteria, for a given game on a PRG, in terms of edge games' values.

It is worth noting that our results can be easily extended to *relaxed hybrid automata* [5], where the strong reset requirement is replaced by a requirement that every cycle in the control graph has a transition that resets all the variables. This extension can be achieved by a refinement of the equivalence relation and a minor modification of the finite graph obtained from it. For clarity, we decided against considering this more general model.

Organisation. Sec. 2 introduces the basic notions used throughout the paper, i.e., definability and decidability, zero-sum games, price-reward graphs, and average-price-per-reward games together with their optimality-equation characterisation. Sec. 3 contains the main technical contribution of the paper: that finite average-price-per-reward games are determined, and that almost optimal controllers exist. In Sec. 4 we state our main results: determinacy and existence of almost-optimal controllers for hybrid average-price-per-reward games.

2 Preliminaries

Here, we introduce key notions that will be used further in the paper, such as definability, decidability, and two-player zero-sum games on price-reward graphs. In Sec. 2.3, we introduce average-price-per-reward games, and optimality equations as means of characterisation (Thm. 4).

Throughout the paper, $\mathbb{R}_\infty$ denotes the set of real numbers augmented with positive and negative infinities, and $\mathbb{R}_+$ and $\mathbb{R}_\oplus$ denote the sets of positive and non-negative reals, respectively. If $G = (V, E)$ is a graph, then for a vertex v, we write vE to denote the set $\{v' : (v, v') \in E\}$ of its successors.

2.1 Definability and decidability

Definability. Let $\mathcal{M} = \langle \mathbb{R}, 0, 1, +, \cdot, \leqslant \rangle$ be the ordered field of reals. We say that a set $X \subseteq \mathbb{R}^n$ is *definable* in $\mathcal{M}$ if it is *first-order definable* in $\mathcal{M}$. The *first-order theory of* $\mathcal{M}$ is the set of all first-order sentences that are true in $\mathcal{M}$. A well-known result by Tarski [17] is that the first-order theory of $\mathcal{M}$ is decidable.

Computability and decidability. For a finite set A, we will say that $(a, x) \in A \times \mathbb{R}^n$ is rational if $x \in \mathbb{Q}^n$. Let $f : X \to \mathbb{R}$ be a partial function, that is defined on a set $D \subseteq X \subseteq \mathbb{R}^n$. We say that f is *approximately computable* if there is an algorithm that for every rational $x \in D$, and every $\varepsilon > 0$, computes a $y \in \mathbb{Q}$ such that $|y - f(x)| < \varepsilon$. It is *decidable* if the following problem is decidable: given a rational $x \in D$ and rational c, decide whether $f(x) \leqslant c$.

We extend the notions of approximate computability, and decidability to functions $f : A \times \mathbb{R}^n \to \mathbb{R}$, where A is finite, by requiring that $f(a, \cdot)$ is respectively: approximately computable, and decidable for every $a \in A$.

Proposition 1. *If a function is decidable then it is approximately computable.*

Proposition 2. *If a real partial function is definable in $\mathcal{M}$ then it is decidable.*

The purpose of the above definitions is to enable us to state conclusions of our definability results. By no means should they be treated as a formalisation of computation over the reals. For models of computing over the reals we refer the reader to [18–20].

2.2 Zero-sum games

In this section we introduce zero-sum games in strategic form, price-per-reward game graphs, and zero-sum price-reward games. Fundamental concepts such as: game value, determinacy, decidability, and optimal strategies are introduced in the context of games in strategic form, and are later lifted to price-reward games. Although our results concern games on price-reward game graphs, the notion of a game in strategic form will be important throughout the paper (for instance, in the formulation of the optimality equations in Sec. 2.3).

Games in strategic form. A zero-sum game is played by two players: Min and Max. Let Σ^{Min}, Σ^{Max} be the sets of *strategies* for players Min and Max respectively. Let $\mathcal{O}$ be the set of outputs, and let $\xi : \Sigma^{\mathrm{Min}} \times \Sigma^{\mathrm{Max}} \to \mathcal{O}$ be a function that, given strategies of players Min and Max, determines the output of the game. Finally let $\mathsf{P} : \mathcal{O} \to \mathbb{R}$, be the payoff function, which given an output determines the payoff. Player Min wants to minimise the payoff, whereas player Max wants to maximise it. A zero-sum game $\eth$ in a *strategic form* is given as $\langle \Sigma^{\mathrm{Min}}, \Sigma^{\mathrm{Max}}, \mathcal{O}, \xi, \mathsf{P} \rangle$. We say that $\eth$ is *definable* if all its components are definable. Recall that definability of a component implicitly implies that it is a subset of $\mathbb{R}^n$.

We define the lower value $\mathsf{Val}_*(\Game) = \sup_{\chi \in \Sigma^{\mathrm{Max}}} \inf_{\mu \in \Sigma^{\mathrm{Min}}} \mathsf{P}(\xi(\mu, \chi))$ and the upper value $\mathsf{Val}^*(\Game) = \inf_{\mu \in \Sigma^{\mathrm{Min}}} \sup_{\chi \in \Sigma^{\mathrm{Max}}} \mathsf{P}(\xi(\mu, \chi))$. Note that $\mathsf{Val}_*(\Game) \leqslant \mathsf{Val}^*(\Game)$, and if these values are equal, then we will refer to them as the value of the game, denoted by $\mathsf{Val}(\Game)$. We will also say that the game is *determined*. Note that, in the definitions above, we allow only pure strategies (i.e., elements of strategy sets).

For all $\mu \in \Sigma^{\mathrm{Min}}$, we define $\mathsf{Val}^{\mu}(\Game) = \sup_{\chi' \in \Sigma^{\mathrm{Max}}} \mathsf{P}(\xi(\mu, \chi'))$. Analogously, for $\chi \in \Sigma^{\mathrm{Max}}$ we define $\mathsf{Val}_{\chi}(\Game) = \inf_{\mu' \in \Sigma^{\mathrm{Min}}} \mathsf{P}(\xi(\mu', \chi))$. For $\varepsilon > 0$, we say that $\mu \in \Sigma^{\mathrm{Min}}$ is *ε-optimal* if we have that $\mathsf{Val}^{\mu}(\Game) \leqslant \mathsf{Val}^*(\Game) + \varepsilon$. We define ε-optimality of strategies for Max analogously.

There are cases in which the desired payoff function is only partially defined on the set of outputs. To remedy this, *lower* $\mathsf{P}_* : \mathcal{O} \to \mathbb{R}$ and *upper* $\mathsf{P}^* : \mathcal{O} \to \mathbb{R}$ payoff functions are used. It is required that $\mathsf{P}_* \leqslant \mathsf{P}^*$. Due to this generalisation, the lower value, and the value of player Max's strategy are defined using the lower payoff, whereas the analogous definitions for the upper value and the value of player Min's strategy use the upper payoff.

Price-reward game graphs. Let $\langle \mathsf{S}, \mathsf{E} \rangle$ be a directed graph and let $\mathsf{S}^{\mathrm{Min}} \uplus \mathsf{S}^{\mathrm{Max}}$ be a partition of S. Let $\mathcal{I}$ be the set of inputs, and let $\Theta^{\mathrm{Min}}, \Theta^{\mathrm{Max}} : \mathsf{E} \to 2^{\mathcal{I}}$ be functions that to every edge, assign the sets of valid inputs. Finally, let $\pi : \mathsf{E} \times \mathcal{I}^2 \to \mathbb{R}$ be a price function, and $\kappa : \mathsf{E} \times \mathcal{I}^2 \to \mathbb{R}_{\oplus}$ be a reward function. A price-reward game graph Γ is given as a tuple $\langle \mathsf{S}^{\mathrm{Min}}, \mathsf{S}^{\mathrm{Max}}, \mathsf{E}, \mathcal{I}, \Theta^{\mathrm{Min}}, \Theta^{\mathrm{Max}}, \pi, \kappa \rangle$. It is said to be *definable* if all its components are definable. When the payoff functions are given, we will refer to Γ as a *price-reward game*.

Intuitively the game is played by moving a token, along the edges, from one state to another. The states are partitioned between players Min and Max. The owner of the state decides along which edge to move the token. The price (reward) of an edge depends on the supplied inputs, one of each is chosen by Min and the other one by Max. The game is played indefinitely. A payoff function determines how the prices (rewards) of individual moves contribute to the overall value of a particular play. The players Min and Max are trying to minimise and maximise (respectively) the value of the payoff function.

We write $s \to_{\theta} s'$ to denote a move, where $e = (s, s') \in \mathsf{E}$ and $\theta \in \Theta^{\mathrm{Min}}(e) \times \Theta^{\mathrm{Max}}(e)$. The price of the move is $\pi(e, \theta)$ and the reward is $\kappa(e, \theta)$. A run is a (possibly infinite) sequence of moves $\rho = s_0 \to_{\theta_1} s_1 \to_{\theta_2} s_2 \cdots$. The set of all valid runs of Γ is denoted by Runs, and its subset of all valid *finite runs* by $\mathrm{Runs}_{\mathrm{fin}}$.

A *state strategy* of player Min is a partial function $\mu^{\mathsf{S}} : \mathrm{Runs}_{\mathrm{fin}} \to \mathsf{E}$ which is defined on all runs ending in $s \in \mathsf{S}^{\mathrm{Min}}$. A strategy is called positional if it can be viewed as a function $\mu^{\mathsf{S}} : \mathsf{S}^{\mathrm{Min}} \to \mathsf{E}$. Given an edge e, an e-strategy of player Min is an element $x \in \Theta^{\mathrm{Min}}(e)$. An *edge strategy* μ^{E} of player Min is a function, that to every edge e assigns an e-strategy.

A strategy μ of player Min is a pair $(\mu^{\mathsf{S}}, \mu^{\mathsf{E}})$ of state and edge strategies. We denote the set of all strategies by Σ^{Min}. We say that μ is positional if μ^{S} is positional. We denote the set of all positional strategies by Π^{Min}. Strategies of player Max are defined analogously.

Given strategies μ and χ of players Min and Max and some state s, we write $\mathsf{Run}(s, \chi, \mu)$ to denote the run starting in s resulting from players playing according to their strategies μ and χ.

Determinacy. Let $\mathsf{P}_* : \text{Runs} \to \mathbb{R}$ and $\mathsf{P}^* : \text{Runs} \to \mathbb{R}$ be the upper and lower payoff functions. Typically, payoff functions are expressions involving prices and rewards of individual transitions.

Given a state s, let $\Game_s = \langle \Sigma^{\text{Min}}, \Sigma^{\text{Max}}, \text{Runs}, \mathsf{Run}(s, \cdot, \cdot), \mathsf{P}^*, \mathsf{P}_* \rangle$. We say that the game Γ is *determined* from s if $\mathsf{Val}_*(\Game_s) = \mathsf{Val}^*(\Game_s)$, and *positionally determined* if $\mathsf{Val}(\Game_s) = \inf_{\mu \in \Pi^{\text{Min}}} \mathsf{Val}^{\mu}(\Game_s) = \sup_{\chi \in \Pi^{\text{Max}}} \mathsf{Val}_{\chi}(\Game_s)$. We say that Γ is determined if it is determined from every state.

For simplicity we will write $\mathsf{Val}(s)$ rather then $\mathsf{Val}(\Game_s)$, in the context of price-reward games, so Val can be viewed as a partial function $\mathsf{S} \to \mathbb{R}$.

Decidability. We will say that a price-reward game Γ is *decidable* if the partial function $\mathsf{Val} : \mathsf{S} \to \mathbb{R}$ is decidable. We emphasise that Val is a partial function because Γ does not have to be determined from every state.

2.3 Average-price-per-reward games

In this section, we introduce average-price-per-reward games, and provide a characterisation of game values using a set of equations, referred to as optimality equations. The key result is Thm. 4, which states that solutions to optimality equations coincide with game values.

The results presented here are general, and will be applied to finite average-price-per-reward games (Sec. 3) as well as to their hybrid counterparts (Sec. 4). The fact that, in both cases, the game values are characterised using optimality equations will be used in the proof of the reduction from hybrid games to finite games (Sec. 4). Notions and arguments similar to those introduced here have been used in the past [6, 14]. We decided to state them in full detail, because they form an important part of our reasoning and provide valuable insight.

The goal of player Min in the *average-price-per-reward game* Γ is to minimise the average price-over-reward ratio in a run, and the goal of player Max is to maximise it. We define the upper and lower payoff functions in the following way:

$$\mathsf{P}^*(\rho) = \limsup_{n \to \infty} \frac{\sum_{i=0}^{n} \pi(e_{i+1}, \theta_{i+1})}{\sum_{i=0}^{n} \kappa(e_{i+1}, \theta_{i+1})} \quad \text{and} \quad \mathsf{P}_*(\rho) = \liminf_{n \to \infty} \frac{\sum_{i=0}^{n} \pi(e_{i+1}, \theta_{i+1})}{\sum_{i=0}^{n} \kappa(e_{i+1}, \theta_{i+1})},$$

where ρ is an infinite run, $s_i \to_{\theta_{i+1}} s_{i+1}$ and $e_i = (s_i, s_{i+1})$ for all $i \geqslant 0$.

To guarantee that the payoffs, as introduced above, are always well-defined we introduce the notions of reward divergence and price-reward boundedness.

We say that Γ is $\Omega(f(n))$*-reward divergent* if, for every run ρ, the function $n \mapsto \sum_{i=0}^{n} \kappa(s_i, \theta_{i+1})$ is in $\Omega(f(n))$. We assume that Γ is $\Omega(n)$-reward divergent. Linear (i.e., $\Omega(n)$) reward divergence is required in the proof of Thm. 4. In the remainder of the paper $c > 0$ will be the largest number such that, for every run ρ, we have $n \mapsto \sum_{i=0}^{n} \kappa(s_i, \theta_{i+1}) \geqslant c \cdot n$.

Additionally, we require that Γ is *price-reward bounded*, i.e., $|\pi| < M$ and $|\kappa| < M$ for some M. This is necessary to assure that edge games, as introduced below, are determined. Moreover, without loss of generality, we assume that the games are non-blocking, i.e., there are no sink states.

The divergence requirement can be seen as a generalisation of the non-zenoness requirement to rewards (as in [8]); we want to prevent runs that admit finite rewards. Note that if the reward is simply time, then we get the non-zenoness condition. Also note that one can can guarantee $\Omega(n)$-reward divergence by claiming that $\kappa > c$ for some $c > 0$.

Optimality equations. Let Γ be a price-per-reward game. For every edge e, we introduce a game $\partial_e(g) = \langle \Theta^{\text{Min}}(e), \Theta^{\text{Max}}(e), \Theta^{\text{Min}}(e) \times \Theta^{\text{Max}}(e), id, \mathsf{P}_e(g) \rangle$, where g is a real-valued parameter, and $\mathsf{P}_e(g) = \pi(e) - \kappa(e) \cdot g$. We will refer to it as an *edge game*. Note that, for every $e \in \mathsf{E}$ and $g \in \mathbb{R}$, we have that $\partial_e(g)$ is determined and definable.

Let $G, B : \mathsf{S} \to \mathbb{R}$ such that the range of G is finite, and B is bounded. We say that a pair of functions (G, B) is a solution of *optimality equations* for Γ, denoted by $(G, B) \models \text{Opt}(\Gamma)$, if the following conditions hold for all states $s \in \mathsf{S}^{\text{Min}}$:

$$G(s) = \min_{(s,s') \in \mathsf{E}} \{G(s')\} \tag{1}$$

$$B(s) = \inf_{(s,s') \in \mathsf{E}} \{\mathsf{Val}(\partial_{(s,s')}(G(s'))) + B(s') \; : \; G(s) = G(s')\} \tag{2}$$

and if analogous two equations hold for all states in S^{Max}, with the only difference that min is substituted by max and inf by sup. The two functions G and B are called *gain* and *bias* respectively.

Remark 3. If Γ is definable then $\text{Opt}(\Gamma)$ is first-order expressible in $\mathcal{M}$.

Theorem 4. *If $(G, B) \models Opt(\Gamma)$ then for every state $s \in \mathsf{S}$, the average-price-per-reward game Γ from s is determined and we have $\mathsf{Val}(s) = G(s)$. Moreover, for every $\varepsilon > 0$, positional ε-optimal strategies exist for both players.*

Corollary 5. *If there exists definable (G, B) such that $(G, B) \models Opt(\Gamma)$ and Γ is definable, then positional ε-optimal strategies are definable.*

The theorem and corollary follow from the following two lemmas and their proofs, which imply that for all states $s \in \mathsf{S}$, we have $\mathsf{Val}^*(s) \leq G(s)$ and $\mathsf{Val}_*(s) \geq G(s)$, respectively.

Lemma 6. *Let $(G, B) \models Opt(\Gamma)$. Then for all $\varepsilon > 0$, there is $\mu_\varepsilon \in \Pi_{Min}$ such that for all $\chi \in \Sigma^{Max}$ and for all $s \in \mathsf{S}$, we have $\mathsf{P}^*(\mathsf{Run}(s, \mu_\varepsilon, \chi)) \leq G(s) + \varepsilon$.*

Lemma 7. *Let $(G, B) \models Opt(\Gamma)$. Then for all $\varepsilon > 0$, there is $\chi_\varepsilon \in \Pi_{Max}$ such that for all $\mu \in \Sigma^{Min}$ and for all $s \in \mathsf{S}$, we have $\mathsf{P}_*(\mathsf{Run}(s, \mu, \chi_\varepsilon)) \geq G(s) - \varepsilon$.*

We omit the proof of Lem. 7 as it is similar to the proof of Lem. 6.

Proof. We prove Lem. 6 by observing that, for every $\varepsilon' > 0$, $g \in \mathbb{R}$, and an edge e, player Min can choose $x^e_{\varepsilon'} \in \Theta^{\mathrm{Min}}(e)$ such that $\mathsf{Val}^{x^e_{\varepsilon'}}(\eth_e(g)) \leqslant \mathsf{Val}(\eth_e(g)) + \varepsilon'$. Moreover, for every state $s \in \mathsf{S}^{\mathrm{Min}}$, player Min can choose an edge $e = (s, s')$ such that:

$$G(s) = G(s')$$
$$B(s) \geqslant \mathsf{P}_e(G(s'))(x^e_{\varepsilon}, y) + B(s') - \varepsilon', \text{ for all } y \in \Theta^{\mathrm{Max}}(e).$$

We will call this choice, of an edge and an edge strategy, ε'-optimal. It remains to show that if, in every $s \in \mathsf{S}^{\mathrm{Min}}$, $\mu_\varepsilon(s)$ is a $(c \cdot \varepsilon)$-optimal choice, then μ_ε is ε-optimal.

Let $\varepsilon > 0$ and $\mu_\varepsilon \in \Pi^{\mathrm{Min}}$ be ε-optimal for every state, and let $\chi \in \Sigma^{\mathrm{Max}}$ be arbitrarily chosen. If $s_i \rightarrow_{\theta_{i+1}} s_{i+1}$ is the $i+1$-th step of $\mathsf{Run}(s, \mu_\varepsilon, \chi)$, then $G(s_i) \geqslant G(s_{i+1})$. The range of G is finite, hence there is $K \in \mathbb{N}$ such that, for all $i \geqslant K$, $G(s_i) = g$, where $g = G(s_K)$.

Let $N \geqslant K$. For $i = K, \ldots, N$, the following holds, $B(s_i) \geqslant \mathsf{P}_e(g)(\theta_{i+1}) + B(s_{i+1}) - c \cdot \varepsilon$. If we sum up the $N - K + 1$ inequalities ($\mathsf{P}_e(g)(\theta) = \pi(e, \theta) - \kappa(e, g) \cdot g$), we get:

$$\sum_{i=K}^{N-1} B(s_i) \geqslant \sum_{i=K+1}^{N} \pi(e_i, \theta_i) - g \cdot \sum_{i=l+1}^{k} \kappa(e_i, \theta_i) + \sum_{i=K+1}^{N} B(s_i) - (N - K + 1) \cdot c \cdot \varepsilon$$

That simplifies to:

$$\frac{B(s_K) - B(s_N)}{\sum_{i=K+1}^{N} \kappa(e_i, \theta_i)} + g \geqslant \frac{\sum_{i=K+1}^{N} \pi(e_i, \theta_i) - (N - K + 1) \cdot c \cdot \varepsilon}{\sum_{i=K+1}^{N} \kappa(e_i, \theta_i)} \geqslant \mathsf{P}^*(\mathsf{Run}(s, \mu_\varepsilon, \chi)) - \varepsilon$$

Recall that B is bounded, and that Γ is $\Omega(n)$-reward divergent with a constant c (which implies that $((N - K + 1) \cdot c \cdot \varepsilon) / \sum_{i=K+1}^{N} \kappa(e_i, \theta_i) \leqslant \varepsilon$). This yields the desired result. □

3 Finite average-price-per-reward games

In this section we state (Thm. 8) and prove (Cor. 13) our technical results, i.e., that finite average-price-per-reward games are determined and decidable[1].

To guarantee uniqueness of the constructions, and for technical convenience, we fix a linear order on the states of the game graph. Given a subgraph $S \subseteq \Gamma$, $\min(S)$ denotes the smallest state in S.

[1] By finite we mean, that the directed graph $\langle \mathsf{S}, \mathsf{E} \rangle$ is finite.

Theorem 8. *Finite average-price-per-reward games are positionally determined and decidable.*

We prove the theorem using the optimality-equation characterisation from Sec. 2.3, and by showing that, in the case of finite price-reward graphs, solutions to optimality equations exist.

Note that we can apply the results from Sec. 2.3 to finite graphs, because gain and bias always have finite ranges.

Strategy subgraphs. Let Γ be a price-reward game graph. Let μ^{S} be a positional state strategy for player Min. Such a strategy induces a subgraph of Γ, where the E relation is substituted by E_μ defined as $E_\mu = \{(s, s') \ : \ s \in \mathsf{S}^{\mathrm{Min}}$ and $\mu^{\mathsf{E}}(s) = s', \text{ or } s \in \mathsf{S}^{\mathrm{Max}}\}$. We denote this game graph by $\Gamma_{\mu^{\mathsf{S}}}$.

A finite connected price-reward game graph of out-degree one is called a *sun*. Such a graph contains a unique cycle, referred to as the *rim*. States which are on the rim are called *rim states* and the remaining ones are called *ray states*.

Remark 9. If $\mu \in \Pi^{\mathrm{Min}}$, $\chi \in \Pi^{\mathrm{Max}}$, and Γ is a price-reward game graph, then $\Gamma_{\mu^{\mathsf{S}}\chi^{\mathsf{S}}}$ is a set of suns.

Game graphs of out-degree one. In price-reward game graphs of out-degree one, strategies of both players are reduced to edge-strategies only. Without loss of generality, we can assume that the price-reward game Γ is defined on a single sun. We now provide a characterisation of upper and lower game values using the values of the rim edge games.

Lemma 10. *Let Γ be a price-reward game defined on a sun, and let $e_1, \ldots, e_k$ denote the edges that form the rim of that sun. Given a parameter $p \in \mathbb{R}$, the following is true for every state s:*

- *If $\sum_{i=1}^{k} \mathsf{Val}(\eth_{e_i}(p)) \geqslant 0$ then $p \leqslant \mathsf{Val}^*(s)$,*
- *If $\sum_{i=1}^{k} \mathsf{Val}(\eth_{e_i}(p)) \leqslant 0$ then $p \geqslant \mathsf{Val}_*(s)$.*

Strict inequalities on the left hand side imply strict inequalities on the right hand side.

Proof. The proof is similar to that of Lemmas 6 and 7. We only sketch the proof of the first statement, as the other is symmetric.

Let χ be a strategy of player Max such that it is $c \cdot \varepsilon$-optimal for every edge game $\eth_{e_i}(p)$, for some $\varepsilon > 0$ and $i = 1, \ldots, k$. If μ is a strategy of player Min, then for every edge e_i:

$$\pi(e_i, \chi(e_i), \mu(e_i)) - \kappa(e_i, \chi(e_i), \mu(e_i)) \cdot p + \varepsilon \geqslant \mathsf{Val}(\eth_{e_i}(p))$$

if we add up the k inequalities we get:

$$\sum_{i=1}^{k} \pi(e_i, \chi(e_i), \mu(e_i)) - \sum_{i=1}^{k} \kappa(e_i, \chi(e_i), \mu(e_i)) \cdot p + k \cdot c \cdot \varepsilon \geqslant 0$$

which gives:

$$\frac{\sum_{i=1}^{k} \pi(e_i, \chi(e_i), \mu(e_i))}{\sum_{i=1}^{k} \kappa(e_i, \chi(e_i), \mu(e_i))} + \varepsilon \geqslant p$$

This, due to the arbitrary choice of ε and μ, finishes the proof. □

Theorem 11. *Solutions to optimality equations for average-price-per-reward games on graphs of out-degree one exist.*

Proof. a finite average-price-per-reward game on a graph of out-degree one, and let S be one of the suns. For every state, both the upper and lower values are finite (recall that Γ is price-reward bounded and linearly reward divergent). Using binary search, together with Lem. 10, it follows that they are indeed equal.

Let g be the value of the game on sun S. We set the gain of all states to g, and the bias of $\min(S)$ to zero. The bias of the remaining states is set to the weight of the shortest path to $\min(S)$, assuming $\mathsf{Val}(\eth_e(g))$ to be the weight on the edge e. Gain and bias functions defined this way satisfy optimality equations. □

General case. We have proved that games on graphs of out-degree one are determined. We will now use this result to prove determinacy in the general case.

Let μ^{S} and χ^{S} be state strategies for players Min and Max respectively, and let (G, B) be gain and bias functions such that $(G, B) \models \mathrm{Opt}(\Gamma_{\mu^{\mathsf{S}}\chi^{\mathsf{S}}})$. Given $s \in \mathsf{S}^{\mathrm{Min}}$ and $e = (s, s') \in \mathsf{E} \setminus \mathsf{E}_{\mu^{\mathsf{S}}\chi^{\mathsf{S}}}$, we say that e is an *improvement* of μ^{S}, with respect to χ^{S}, if $G(s) > G(s')$, or $G(s) = G(s')$ and $B(s) > \mathsf{Val}(\eth_e(G(s)) + B(s')$. A strategy μ'^{S} is an improvement of μ^{S} with respect to χ^{S} if for every state s, either $\mu^{\mathsf{S}}(s) = \mu'^{\mathsf{S}}(s)$, or $\mu'^{\mathsf{S}}(s) = s'$ and (s, s') is an improvement of μ^{S} with respect to χ^{S}. An improvement is strict if $\mu^{\mathsf{S}} \neq \mu'^{\mathsf{S}}$. An improvement of χ^{S} is defined similarly.

We say that χ^{S}, a state strategy for player Max, is a *best response* to μ^{S}, a state strategy of player Min, if there are no possible improvements of χ^{S} with respect to μ^{S}.

To prove the existence of best response strategies we apply Thm. 12 and the fact that the set of edge strategies is finite, to average-price-per-reward games, in which all the states belong to only one player.

Theorem 12. *Let μ^{S} be a state strategy of player Min, χ^{S} a best response strategy of player Max, and (G, B) gain and bias such that $(G, B) \models Opt(\Gamma_{\mu^{\mathsf{S}}\chi^{\mathsf{S}}})$. If μ'^{S} is an improvement of μ^{S} with respect to χ^{S}, χ'^{S} is a best response to μ'^{S}, and $(G', B') \models Opt(\Gamma_{\mu'^{\mathsf{S}}\chi'^{\mathsf{S}}})$, then the following holds for every state s:*

1. *$G(s) < G'(s)$, or*
2. *$G(s) = G'(s)$ and $B(s) \leqslant B'(s)$.*

Moreover, if $\mu^{\mathsf{S}} \neq \mu'^{\mathsf{S}}$ then $(G, B) \neq (G', B')$.

Proof. Consider the game graph $\Gamma_{\mu'^{\mathsf{S}}\chi'^{\mathsf{S}}}$. For every edge $e = (s, s')$, either i) $G(s) > G(s')$, or ii) $G(s) = G(s')$ and $B(s) \geqslant \mathsf{Val}(\eth_e(G(s))) + B(s')$.

Using the same argument as in Lem. 6, we show that $G \geqslant G'$ for all cycles in $\Gamma_{\mu'^{\mathsf{s}}\chi'^{\mathsf{s}}}$, and that $G > G'$ for cycles that did not exist in $\Gamma_{\mu^{\mathsf{s}}\chi^{\mathsf{s}}}$. This proves (1).

Let s be a vertex such that $G(s) = G'(s)$, and let S be a sun in $\Gamma_{\mu'^{\mathsf{s}}\chi'^{\mathsf{s}}}$ such that, $s \in S$. If $s_0, \ldots, s_k$ is the path from s to $\min(S)$ then, for every (s_i, s_{i+1}), $B(s_i) \geqslant \mathsf{Val}(\eth_{(s_i,s_{i+1})}(G(s))) + B(s_{i+1})$. If we sum up, the k inequalities, we get that $B(s)$ is no less then the weight of $s_0, \ldots, s_k$, assuming $\mathsf{Val}(\eth_{(s_i,s_{i+1})}(G(s)))$ to be the weight of edge (s_i, s_{i+1}), which in turn is equal to $G'(s)$. □

Corollary 13. *Solutions to optimality equations for average-price-per-reward games exist.*

Proof. The set of edge strategies for both players is finite. This, together with Thm. 12, guarantees the existence of mutual best response edge strategies. The rest follows from Thm. 11. □

Theorem 14. *Finite definable average-price-per-reward games are decidable.*

Proof. $\mathrm{Opt}(\Gamma)$ is finite hence (G, B) such that $(G, B) \models \mathrm{Opt}(\Gamma)$, is definable (by Rem. 3). □

4 Games on hybrid automata with strong resets

We introduce hybrid automata with strong resets and define price-reward hybrid games on these automata. The main result is that the hybrid average-price-per-reward games are determined and decidable (Thm. 16). To obtain the result, we reduce hybrid average-price-per-reward games to finite average-price-per-reward games.

Our definition of a hybrid automaton varies from that used in [12, 13], as we hide the dynamics of the system into guard functions. This approach allows for cleaner and more succinct notation and exposition, without loss of generality [6].

Price-reward hybrid automata with strong resets. Let L be a finite set of *locations*. Fix $n \in \mathbb{N}$ and define the set of *states* $\mathsf{S} = \mathsf{L} \times \mathbb{R}^n$. Let A be a finite set of *actions*, and define the set of *times* $\mathsf{T} = \mathbb{R}_{\oplus}$. We refer to action-time pairs $(a, t) \in \mathsf{A} \times \mathsf{T}$ as *timed actions*. A *price-reward hybrid automaton with strong resets (PRHASR)* $\mathcal{H} = \langle \mathsf{L}, \mathsf{A}, \mathsf{G}, \mathsf{R}, \pi, \kappa \rangle$ consists of finite sets L of *locations* and A of *actions*, a *guard* function $\mathsf{G} : \mathsf{A} \to 2^{\mathsf{S} \times \mathsf{T}}$, a *reset* function $\mathsf{R} : \mathsf{A} \to 2^{\mathsf{S}}$, a continuous *price* function $\pi : \mathsf{S} \times (\mathsf{A} \times \mathsf{T}) \to \mathbb{R}$, and a continuous *reward* function $\kappa : \mathsf{S} \times (\mathsf{A} \times \mathsf{T}) \to \mathbb{R}_{\oplus}$. We say that $\mathcal{H}$ is a *definable PRHASR* if the functions $\mathsf{G}, \mathsf{R}, \pi$, and κ are definable.

For states $s, s' \in \mathsf{S}$ and a timed action $(a, t) \in \mathsf{A} \times \mathsf{T}$, we write $s \xrightarrow{a}_t s'$ iff $(s, t) \in \mathsf{G}(a)$ and $s' \in \mathsf{R}(a)$. If $s, s' \in \mathsf{S}$, $\tau = (a, t) \in \mathsf{A} \times \mathsf{T}$, and $s \xrightarrow{a}_t s'$ then we write $s \xrightarrow{\tau} s'$. We define the *move* function $\mathsf{M} : \mathsf{S} \to 2^{\mathsf{A} \times \mathsf{T}}$ by $\mathsf{M}(s) = \{(a, t) : (s, t) \in \mathsf{G}(a)\}$. Note that M is definable if G is definable. A *run* from state $s \in \mathsf{S}$ is a sequence $\langle s_0, \tau_1, s_1, \tau_2, s_2, \ldots \rangle \in \mathsf{S} \times ((\mathsf{A} \times \mathsf{T}) \times \mathsf{S})^\omega$ such that $s_0 = s$, and for all $i \geq 0$, we have $s_i \xrightarrow{\tau_{i+1}} s_{i+1}$.

Hybrid games with strong resets. A *hybrid game with strong resets (HGSR)* $\Gamma = \langle \mathcal{H}, \mathsf{M}^{\mathrm{Min}}, \mathsf{M}^{\mathrm{Max}} \rangle$ consists of a PRHASR $\mathcal{H} = \langle \mathsf{L}, \mathsf{A}, \mathsf{G}, \mathsf{R}, \pi, \kappa \rangle$, a *Min-move* function $\mathsf{M}^{\mathrm{Min}} : \mathsf{S} \to 2^{\mathsf{A} \times \mathsf{T}}$ and a *Max-move* function $\mathsf{M}^{\mathrm{Max}} : \mathsf{S} \times (\mathsf{A} \times \mathsf{T}) \to 2^{\mathsf{A} \times \mathsf{T}}$. We require that for all $s \in \mathsf{S}$, we have $\mathsf{M}^{\mathrm{Min}}(s) \subseteq \mathsf{M}(s)$, and that for all $\tau \in \mathsf{M}^{\mathrm{Min}}(s)$, we have $\mathsf{M}^{\mathrm{Max}}(s, \tau) \subseteq \mathsf{M}(s)$. W.l.o.g., we assume that for all $s \in \mathsf{S}$, we have $\mathsf{M}^{\mathrm{Min}}(s) \neq \emptyset$, and that for all $\tau \in \mathsf{M}^{\mathrm{Min}}(s)$, we have $\mathsf{M}^{\mathrm{Max}}(s, \tau) \neq \emptyset$. If $\mathcal{H}$ and the move functions are definable, then we say that Γ is *definable*.

In the reminder of the paper, we consider price-reward HGSRs. For simplicity, we refer to them as hybrid price-reward games or, when the price-reward aspect is irrelevant, just hybrid games.

A hybrid game is played in rounds. In every round, the following three steps are performed by the two players Min and Max from the current state $s \in \mathsf{S}$.

1. Player Min proposes a timed action $\tau \in \mathsf{M}^{\mathrm{Min}}(s)$.
2. Player Max responds by choosing a timed action $\tau' = (a', t') \in \mathsf{M}^{\mathrm{Max}}(s, \tau)$. This choice determines the price and reward contribution of the round ($\pi(s, \tau')$ and $\kappa(s, \tau')$ respectively).
3. Player Max chooses a state $s' \in \mathsf{R}(a')$, i.e., such that $s \xrightarrow{\tau'} s'$. The state s' becomes the current state for the next round.

A *play* of the game Γ from state s is a sequence $\langle s_0, \tau_1, \tau_1', s_1, \tau_2, \tau_2', s_2, \ldots \rangle \in \mathsf{S} \times ((\mathsf{A} \times \mathsf{T}) \times (\mathsf{A} \times \mathsf{T}) \times \mathsf{S})^\omega$, such that $s_0 = s$, and for all $i \geq 0$, we have $\tau_{i+1} \in \mathsf{M}^{\mathrm{Min}}(s_i)$ and $\tau_{i+1}' \in \mathsf{M}^{\mathrm{Max}}(s_i, \tau_{i+1})$. Note that if $\langle s_0, \tau_1, \tau_1', s_1, \tau_2, \tau_2', s_2, \ldots \rangle$ is a play then the sequence $\langle s_0, \tau_1', s_1, \tau_2', s_2, \ldots \rangle$ is a run of the hybrid automaton $\mathcal{H}$.

A hybrid game with strong resets can be viewed as a game on an infinite price-reward game graph, with fixed costs and rewards assigned to edges. The set of states S' is a subset of: $\mathsf{S} \cup (\mathsf{S} \times (\mathsf{A} \times \mathsf{T})) \cup ((\mathsf{A} \times \mathsf{T}))$. The E' relation is defined as follows: $(s, (s, \tau)) \in \mathsf{E}'$ iff $\tau \in \mathsf{M}^{\mathrm{Min}}(s)$, and $((s, \tau), \tau') \in \mathsf{E}$ iff $\tau' \in \mathsf{M}^{\mathrm{Max}}(s, \tau)$, and $((a', t'), s') \in \mathsf{E}'$ iff $s' \in \mathsf{R}(a')$.

We define $\Gamma' = \langle \mathsf{S}, \mathsf{S}' \setminus \mathsf{S}, \mathsf{E}', \pi', \kappa' \rangle$, where for an edge $e = ((s, \tau), (a', t'))$, we set $\pi'(e) = \pi(s, t')$ and $\kappa'(e) = \kappa(s, t')$, and for all other edges we set them to 0. Additionally, we require that $\mathsf{S}' \setminus \mathsf{S}$ contains all states reachable from S and does not contain those that are not. In the defintion of Γ', we omitted the inputs, as neither the prices nor the rewards depend on them.

Remark 15. For all $(a, t), (a', t') \in \mathsf{S}'$, if $a = a'$ then $(a, t)E' = (a', t')E'$. This is a consequence of the strong reset property of $\mathcal{H}$.

It is clear that plays of Γ directly correspond to runs on Γ'. Moreover, any run of Γ' uniquely determines a run of $\mathcal{H}$. We will use this fact to, lift the concepts introduced for price-reward games to hybrid price-reward games. We will say that the hybrid game Γ has a property P if Γ' has this property.

Hybrid average-price-per-reward games. In the following, we lift the concept of average-price-per-reward games, as defined in Sec. 2.3, to hybrid price-reward games. We state and prove the main result of the paper:

Theorem 16. *Average-price-per-reward hybrid games are positionally determined and decidable.*

We prove the theorem through a reduction to finite average-price-per-reward games. To obtain the corresponding finite price-reward graph we use an equivalence relation on the state space of the hybrid automaton.

We define the lower and upper payoffs as follows. For a run $\rho = \langle s_0, \tau_1 s_1, \tau_2 \ldots \rangle$ of $\mathcal{H}$, we define the lower payoff P_* and the upper payoff P^* by

$$\mathsf{P}_*(\rho) = \liminf_{n\to\infty} \frac{\sum_{i=0}^{n-1} \pi(s_i, \tau_{i+1})}{\sum_{i=0}^{n-1} \kappa(s_i, \tau_{i+1})} \qquad \mathsf{P}^*(\rho) = \limsup_{n\to\infty} \frac{\sum_{i=0}^{n-1} \pi(s_i, \tau_{i+1})}{\sum_{i=0}^{n-1} \kappa(s_i, \tau_{i+1})}$$

Note that these payoffs are exactly the same, as the average-price-per-reward payoffs for runs starting in $\mathsf{S} \subseteq \mathsf{S}'$ in Γ' (we therefore require that Γ is $\Omega(n)$-divergent and price convergent). This enables us to use the optimality equation characterisation and results from Sec. 2.3. Using Rem. 15 and the fact that A is a finite set, we guarantee that gain has a finite range, and that bias is bounded.

We will also say that $\mathrm{Opt}(\Gamma')$ is the set of optimality equations for the hybrid game Γ, denoted by $\mathrm{Opt}(\Gamma)$. Let $G, B : \mathsf{S} \cup \big(\mathsf{S} \times (\mathsf{A} \times \mathsf{T})\big) \cup \mathsf{A} \to \mathbb{R}$. The optimality equations for Γ' take the following form: if $s \in \mathsf{S}$, then

$$G(s) = \min_{\tau \in \mathsf{M}^{\mathrm{Min}}(s)} \{G(s,\tau)\}, \tag{3}$$

$$B(s) = \inf_{\tau \in \mathsf{M}^{\mathrm{Min}}(s)} \{B(s,\tau) \ : \ G(s,\tau) = G(s)\}; \tag{4}$$

if $s \in \mathsf{S}$ and $\tau \in \mathsf{M}^{\mathrm{Min}}(s)$, then

$$G(s,\tau) = \max_{(a',t') \in \mathsf{M}^{\mathrm{Max}}(s,\tau)} \{G(a')\}, \tag{5}$$

$$B(s,\tau) = \sup_{(a',t') \in \mathsf{M}^{\mathrm{Max}}(s,\tau)} \{\pi(s,a',t') - \kappa(s,a',t') \cdot G(a') + B(a') \ : \ G(a') = G(s,\tau)\}; \tag{6}$$

and if $a \in \mathsf{A}$

$$G(a) = \max_{s \in \mathsf{R}(a)} \{G(s)\}, \qquad B(a) = \sup_{s \in \mathsf{R}(a)} \{B(s) \ : \ G(s) = G(a)\}.$$

The last pair of equations is a generic pair of equations for all states $(a,t) \in \mathsf{S}'$. This is valid by Rem. 15. We have written the equations taking into account the fixed price and rewards in Γ'.

Solving hybrid average-price-per-reward games. We show that hybrid average-price-per-reward games are determined and decidable.

In order to establish our results, we use an equivalence relation over the state space of the hybrid game Γ, as introduced in [6]. This relation is of finite index, and its equivalence classes are used to construct a finite price-reward game graph $\widehat{\Gamma}$.

We characterise the game values using optimality equations from Sec. 2.3, and prove that solutions to $\mathrm{Opt}(\widehat{\Gamma})$ coincide with the solutions to $\mathrm{Opt}(\Gamma)$. This, together with the results from Sec. 3 proves that hybrid average-price-per-reward games are determined.

Recall the definition of equivalence relation $\sim$, and the details of the finite graph construction from [6]. We obtain the finite price-reward game graph $\widehat{\Gamma} = (\widehat{\mathsf{S}}^{\mathrm{Min}}, \widehat{\mathsf{S}}^{\mathrm{Max}}, \widehat{\mathsf{E}}, \widehat{\mathcal{I}}, \widehat{\Theta}^{\mathrm{Min}}, \widehat{\Theta}^{\mathrm{Max}}, \widehat{\pi}, \widehat{\kappa})$ from $\Gamma = (\mathcal{H}, \mathsf{M}^{\mathrm{Min}}, \mathsf{M}^{\mathrm{Max}})$ the following way. The finite graph $(\widehat{\mathsf{S}}, \widehat{\mathsf{E}})$ is given by:

$$
\begin{aligned}
\widehat{\mathsf{S}} &= \mathsf{A} \cup \mathsf{S}/\!\sim\, \cup \{(Q, a, A') \,:\, Q \in \mathsf{S}/\!\sim \text{ and } (a, A') \in \mathsf{A}^{\mathrm{MinMax}}(Q, \mathsf{T})\},\\
\widehat{\mathsf{E}} &= \{(a, Q) \,:\, Q \subseteq \mathsf{R}(a)\} \cup \{\big(Q, (Q, a, A')\big) \,:\, (a, A') \in \mathsf{A}^{\mathrm{MinMax}}(Q, \mathsf{T})\}\\
&\qquad \cup \{\big((Q, a, A'), a'\big) \,:\, a' \in A'\},
\end{aligned}
$$

and the partition of $\widehat{S}$ is given by $\widehat{\mathsf{S}}^{\mathrm{Min}} = \mathsf{S}/\!\sim$ and $\widehat{\mathsf{S}}^{\mathrm{Max}} = \widehat{\mathsf{S}} \setminus \widehat{\mathsf{S}}^{\mathrm{Min}}$. The set of inputs is $\widehat{\mathcal{I}} = \{\ni\} \cup [\mathsf{S} \to \mathsf{A} \times \mathsf{T}] \cup \mathsf{S} \times [\mathsf{A} \times \mathsf{T} \to \mathsf{A} \times \mathsf{T}]$ ($\ni$ serves as a special input for edges that will bear a fixed 0 price and reward). For an edge $e = (Q, a, A')$, let $\widehat{\Theta}^{\mathrm{Min}}(e) \subseteq [Q \to \mathsf{A} \times \mathsf{T}]$ be such that for every $s \in Q$ and $f \in \widehat{\Theta}^{\mathrm{Min}}(e)$, we have that $f(s) \in \mathsf{M}^{\mathrm{Min}}(s)$, and let $\widehat{\Theta}^{\mathrm{Max}}(e) \subseteq Q \times [\mathsf{A} \times \mathsf{T} \to \mathsf{A} \times \mathsf{T}]$ be such that for every $s \in Q$, $\tau \in \mathsf{M}^{\mathrm{Min}}(s)$ and $(s, f) \in \widehat{\Theta}^{\mathrm{Max}}(e)$, we have that $f(\tau) \in \mathsf{M}^{\mathrm{Max}}(s, \tau)$. Let $f \in \widehat{\Theta}^{\mathrm{Min}}(e)$ and $(s, f') \in \widehat{\Theta}^{\mathrm{Max}}(e)$, we define the price (reward) of that edge as $\widehat{\pi}(e)(f, (s, f')) = \pi(s, f'(f(s)))$ ($\widehat{\kappa}(e)(f, (s, f')) = \kappa(s, f'(f(s)))$). For the remaining edges we set $\widehat{\Theta}^{\mathrm{Min}}$ and $\widehat{\Theta}^{\mathrm{Max}}$ to $\{\ni\}$, and their price (reward) to 0.

Theorem 17. *Let Γ be a hybrid average-price-per-reward game and let $(\widehat{G}, \widehat{B}) \models$ Opt($\widehat{\Gamma}$). If $G, B : \mathsf{S} \cup \big(\mathsf{S} \times (\mathsf{A} \times \mathsf{T})\big) \cup \mathsf{A} \to \mathbb{R}$ are such that $G(a) = \widehat{G}(a)$ and $B(a) = \widehat{B}(a)$ for all $a \in \mathsf{A}$, and satisfy equations (3–6), then $(G, B) \models$ Opt(Γ).*

Corollary 18. *Definable average-price-per-reward hybrid games with strong resets are decidable.*

References

1. Alur, R., Courcoubetis, C., Henzinger, T.A., Ho, P.H., Nicollin, X., Olivero, A., Sifakis, J., Yovine, S.: The algorithmic analysis of hybrid systems. Theoretical Computer Science **138** (1995) 3–34
2. Alur, R., Dill, D.: A theory of timed automata. Theoretical Computer Science **126** (1994) 183–235
3. Henzinger, T.A.: The theory of hybrid automata. In: LICS 1996, IEEE Computer Society Press (1996) 278–292
4. Lafferriere, G., Pappas, G.J., Sastry, S.: O-minimal hybrid systems. Mathematics of Control, Signals, and Systems **13** (2000) 1–21
5. Gentilini, R.: Reachability problems on extended o-minimal hybrid automata. In: FORMATS'05. LNCS, Springer (2005) 162–176
6. Bouyer, P., Brihaye, T., Jurdziński, M., Lazić, R., Rutkowski, M.: Average-price and reachability-price games on hybrid automata with strong resets. In: FORMATS'08. Volume 5215 of LNCS., Springer (2008) 63–77

7. Brihaye, T., Michaux, C.: On the expressiveness and decidability of o-minimal hybrid systems. Journal of Complexity **21** (2005) 447–478
8. Bouyer, P., Brinksma, E., Larsen, K.G.: Optimal infinite scheduling for multi-priced timed automata. Formal Methods in System Design **32** (2008) 2–23
9. Abdeddam, Y., Asarin, E., Maler, O.: Scheduling with timed automata. Theoretical Computer Science (2006) 272–300
10. Church, A.: Logic, arithmetic and automata. In: Proceedings of the International Congress of Mathematicians. (1962) 23–35
11. Thomas, W.: On the synthesis of strategies in infinite games. In: STACS'95. Volume 900 of LNCS., Springer (2005) 1–13
12. Bouyer, P., Brihaye, T., Chevalier, F.: Control in o-minimal hybrid systems. In: LICS'06, LNCS (2006) 367–378
13. Bouyer, P., Brihaye, T., Chevalier, F.: Weighted o-minimal hybrid systems are more decidable than weighted timed automata! In: LFCS'07. Volume 4514 of LNCS., Springer (2007) 69–83
14. Jurdziński, M., Trivedi, A.: Reachability-time games on timed automata. In: ICALP'07. Volume 4596 of LNCS., Springer (2007) 838–849
15. Brihaye, T., Michaux, C., Rivière, C., Troestler, C.: On o-minimal hybrid systems. In: HSCC'04. Volume 2993 of LNCS., Springer (2004) 219–233
16. Brihaye, T.: A note on the undecidability of the reachability problem for o-minimal dynamical systems. Mathematical Logic Quarterly **52** (2006) 165–170
17. Tarski, A.: A Decision Method for Elementary Algebra and Geometry. University of California Press (1951)
18. Weihrauch, K.: Computable Analysis. Sringer (2000)
19. Meer, K., Michaux, C.: A survey on real structural complexity theory. Bulletin of the Belgian Mathematical Society. **4** (1997) 113–148 Journées Montoises (Mons, 1994).
20. Blum, L., Shub, M., Smale, S.: On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. American Mathematical Society. Bulletin. New Series **21** (1989) 1–46