Admissibility in Quantitative Graph Games*

Romain Brenguier[1], Guillermo A. Pérez[†2], Jean-François Raskin[3], and Ocan Sankur[4]

1 University of Oxford, Oxford, UK
2 Université Libre de Bruxelles (ULB), Brussels, Belgium
3 Université Libre de Bruxelles (ULB), Brussels, Belgium
4 CNRS, Irisa, Rennes, France

Abstract

Admissibility has been studied for games of infinite duration with Boolean objectives. We extend here this study to games of infinite duration with *quantitative* objectives. First, we show that, under the assumption that optimal worst-case and cooperative strategies exist, admissible strategies are guaranteed to exist. Second, we give a characterization of admissible strategies using the notion of adversarial and cooperative values of a history, and we characterize the set of outcomes that are compatible with admissible strategies. Finally, we show how these characterizations can be used to design algorithms to decide relevant verification and synthesis problems.

1998 ACM Subject Classification F.1.1 Automata, D.2.4 Formal methods

Keywords and phrases Quantitative games, Verification, Reactive synthesis, Admissibility

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2016.42

1 Introduction

Two-player *zero-sum* graph games are the most studied mathematical model to formalize the reactive synthesis problem [15, 16]. Unfortunately, this mathematical model is often an abstraction that is too coarse. Realistic systems are usually made up of several components, all of them with their *own* objectives. These objectives are not necessarily antagonistic. Hence, the setting of *non-zero sum* graph games is now investigated in order to *unleash* the full potential of automatic synthesis algorithms for reactive systems, see *e.g.* [9, 2, 5, 6, 14, 12].

For a player with objective φ, a strategy σ is said to be *dominated* by a strategy σ' if σ' does as well as σ with respect to φ against all the strategies of the other players and strictly better for some of them. A strategy σ is *admissible* for a player if it is *not* dominated by any other of his strategies. Clearly, playing a strategy which is not admissible is sub-optimal and a *rational* player should only play admissible strategies. The elimination of dominated strategies can be *iterated* if one assumes that each player knows the other players know that only admissible strategies are played, and so on.

While admissibility is a classical notion for finite games in normal form, see *e.g.* [13] and pointers therein, its generalization to infinite duration games is challenging and was only considered more recently. In 2007, Berwanger was the first to show [2] that admissibility, *i.e.* the avoidance of dominated strategies, is well-behaved in infinite duration n-player non-zero sum turn-based games with perfect information and Boolean outcomes (two possible payoffs:

* This work was partially supported by the ERC Starting Grant inVEST (279499) and EPSRC grant EP/M023656/1.
† Author supported by F.R.S.-FNRS fellowship.

36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2016).
Editors: Akash Lal, S. Akshay, Saket Saurabh, and Sandeep Sen; Article No. 42; pp. 42:1–42:14
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

win or lose). This framework encompasses games with omega-regular objectives. The main contributions of Berwanger were to show that

(i) in all n-player game structures, for all objectives, players have *admissible strategies*, (Berwanger even shows the existence of strategies that survive the iterated elimination of strategies)
(ii) every strategy that is dominated by a strategy is dominated by an admissible strategy, and
(iii) for finite game structures, the set of admissible strategies forms a regular set.

While the iterated admissibility formalizes a strong notion of rationality [1], it has been shown recently that the non-iterated version is strong enough to synthesize relevant strategies for non-zero sum games of infinite duration modelling reactive systems. In [11], Faella considers games played on finite graphs and focuses on the states from which one designated player cannot force a win. He compares several criteria for establishing what is the preferable behavior of this player from those states, eventually settling on the notion of admissible strategy. In [4], starting from the notion of admissible strategy, we have defined a novel rule for the compositional synthesis of reactive systems, applicable to systems made of n components which have each their own objective. We have shown that this synthesis rule leads to solutions which are robust and resilient.

Here, we study the notion of admissible strategy in infinite horizon n-player turn-based *quantitative* games played on a finite game structure. We give a comprehensive picture of the properties related to the existence of such strategies and to their characterization. Contrary to the Boolean case, the number of payoffs in our setting is potentially infinite making the characterization challenging. As in [2], we assume all players have perfect information.

Main contributions. First, contrary to the Boolean case, we show that in the quantitative setting, there are dominated strategies that are not dominated by any admissible strategy (Example 9). Second, we show that the existence of worst-case optimal and cooperatively optimal strategies for all players is a sufficient condition for the existence of admissible strategies (Thm. 4). Additionally, we show that there are games without worst-case optimal or without cooperative optimal strategies that do not have admissible strategies (Lem. 3). Third, we provide a characterization of admissible strategies in terms of antagonistic and cooperative values – that are classical values defined for quantitative games – (Thm. 11) and a characterization of the outcomes compatible with admissible strategies (Thm. 13). While the first characterization allows one to precisely describe admissible strategies, the characterization of the set of outcomes is given in linear temporal logic, and is a useful tool to reason about the outcomes that can be generated by such strategies. Finally, we show how to use the aforementioned characterizations to obtain algorithms to solve relevant decision problems for games with classical quantitative measures such as Inf, Sup, LimInf, LimSup and mean-payoff (Thms. 17, 18, and 19).

Example. Let us consider the game from Fig. 1 to illustrate several notions and decision problems introduced and solved in this paper. The game is played by two players: Player 1, who owns the square vertices, and Player 2, owner of the round vertices. The measure that we consider here is the mean-payoff. (But note that, the arguments we will develop in this example are applicable to the limit inferior and limit superior measures as well.)

First, we note that the (best) worst-case value (or, the antagonistic value) that Player 1 can force is equal to 1, while the antagonistic value for Player 2 is equal to 0. The latter values are meaningful under the hypothesis that the other player is playing fully antagonistically

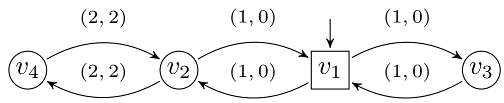

Figure 1 Player 1 controls the square vertices, and Player 2 the round vertices. The payoff of Player i is the mean-payoff of the dimension i of the weights seen along the run.

and not pursuing their own objective. Now, if we account for the fact that Player 2 aims at maximizing his own payoff and so plays only admissible strategies towards this goal, then we conclude that he will never play the edge (v_2, v_1). This is because, from vertex v_2, Player 2 has a strategy to enforce value 2 and taking edge (v_2, v_1) is unreasonable because, in the worst case, from v_1 he will only obtain 0. As we show in Sec. 6, this kind of reasoning can be made formal and automated. We will show that, for games with classical quantitative measures, it can indeed be decided algorithmically if a finite memory strategy given, for instance, as a finite state Moore machine, is admissible or not.

Second, a similar but more subtle reasoning to the one presented above allows us to conclude that Player 1 will eventually play the edge (v_1, v_2). Indeed, from vertex v_1, Player 1 can force a payoff equal to 1 by either taking edge (v_1, v_3) or (v_1, v_2). Nevertheless, it is not reasonable for him to play edge (v_1, v_3) because, while this choice enforces a worst-case payoff equal to 1 (the antagonistic value), playing edge (v_1, v_2) is better because it ensures the same worst-case payoff and additionally leaves a possibility for Player 2 to help him by taking the cycle v_2–v_4, giving him a payoff of 2. If we take into account that the adversary is playing admissible strategies, then, in the words of [4], we can solve the assume-admissible synthesis problem. In this example, we conclude that Player 1 has a strategy to enforce a payoff of 2 against all admissible strategies of Player 2. A strategy which eventually chooses edge (v_1, v_2) ensures this payoff. The formalization of this reasoning and elements necessary for its automation are presented in Sec. 6.

Structure of the paper. Sec. 2 contains definitions. In Sec. 3, we study conditions under which the existence of admissible strategies is guaranteed. In Sec. 4, we give a characterization of admissible strategies, and in Sec. 5, a description of the set of outcomes compatible with admissible strategies. In Sec. 6, we apply our results to solve relevant decision problems on games with classical quantitative measures.

2 Preliminaries

We denote by $\mathbb{R}$ the set of *real numbers*, $\mathbb{Q}$ the set of *rational numbers*, $\mathbb{N}$ the set of *natural numbers*, and $\mathbb{N}_{>0}$ the set of *positive integers*.

A *game* is a tuple $\mathcal{G} = \langle P, (V_i)_{i \in P}, v_{\mathsf{init}}, E, (\mathsf{payoff}_i)_{i \in P} \rangle$ where:

(i) P is the non-empty and finite set of players.

(ii) $V \stackrel{\text{def}}{=} \biguplus_{i \in P} V_i$ where for every $i \in P$, V_i is the finite set of player i's vertices, and $v_{\mathsf{init}} \in V$ is the *initial vertex*.

(iii) $E \subseteq V \times V$ is the set of edges (it is assumed, w.l.o.g., that each vertex in V has at least one outgoing edge.)

(iv) For every i in P, payoff_i is a *payoff function* from infinite paths in the digraph $\langle V, E \rangle$ to $\mathbb{R}$ that, intuitively, player i will attempt to maximize.

An *outcome* ρ is an infinite path in the digraph $\langle V, E \rangle$, *i.e.* an infinite sequence of vertices $(\rho_j)_{j \in \mathbb{N}_{>0}}$ such that $(\rho_j, \rho_{j+1}) \in E$, for all $j \in \mathbb{N}_{>0}$. A finite prefix of an outcome is called a *history.* The *length* $|h|$ of a history $h = (\rho_j)_{1 \le j \le n}$ is n. Given an outcome $\rho = (\rho_j)_{j \in \mathbb{N}_{>0}}$ and an integer k, we write $\rho_{\le k}$ for the history $(\rho_j)_{1 \le j \le k}$, that is, the prefix of length k of ρ. For a history h and a history or outcome ρ, we write $h \subseteq_{\mathsf{pref}} \rho$ if h is a prefix of ρ. If $h \subseteq_{\mathsf{pref}} \rho$, we write $h^{-1} \cdot \rho$ for the unique history (resp. outcome) that satisfies $\rho = h \cdot (h^{-1} \cdot \rho)$. The *first* (resp. *last*) vertex of a history h is $\mathrm{first}(h) = h_1$ (resp. $\mathrm{last}(h) \stackrel{\mathsf{def}}{=} h_{|h|}$). The *longest common prefix* of two outcomes or histories ρ, ρ' is denoted $\mathrm{lcp}(\rho, \rho')$. Given vertex v from $\mathcal{G}$, let us denote the set of *successors of* v by $E_v \stackrel{\mathsf{def}}{=} \{v' \in V \mid (v, v') \in E\}$.

A *strategy* of player i is a function σ_i that maps any history h such that $\mathrm{last}(h) \in V_i$ to a vertex from $E_{\mathrm{last}(h)}$. A *strategy profile* for the set of players $P' \subseteq P$ is a tuple of strategies, one for each player of P'.

Let $\Sigma_i(\mathcal{G})$ be the set of all strategies of player i in $\mathcal{G}$. We write $\Sigma(\mathcal{G}) \stackrel{\mathsf{def}}{=} \prod_{i \in P} \Sigma_i(\mathcal{G})$ for the set of all strategy profiles for P in $\mathcal{G}$, and $\Sigma_{-i}(\mathcal{G})$ for the set of strategy profiles for all players but i in $\mathcal{G}$. We omit $\mathcal{G}$ when it is clear from the context. Given $\sigma_i \in \Sigma_i$ and $\sigma_{-i} = (\sigma_j)_{j \in P \setminus \{i\}} \in \Sigma_{-i}$, we write (σ_i, σ_{-i}) for $(\sigma_j)_{j \in P}$.

A strategy profile $\sigma_P \in \Sigma$ defines a unique *outcome* from any given history h. Formally, $\mathbf{Out}_h(\mathcal{G}, \sigma_P)$ is the outcome $\rho = (\rho_j)_{j \in \mathbb{N}_{>0}}$ such that $\rho_{\le |h|} = h$ and for $j > |h|$, if $\rho_j \in V_i$, then $\rho_{j+1} = \sigma_i(\rho_{\le j})$. Notice that when h is a vertex, then this corresponds to starting the game at that vertex. When $\mathcal{G}$ is clear from the context we shall omit it and write simply $\mathbf{Out}_h(\sigma_P)$. If S_i is a set of strategies for player i, we write $\mathbf{Out}_h(S_i)$ for $\{\rho \mid \exists \sigma_i \in S_i, \sigma_{-i} \in \Sigma_{-i} : \mathbf{Out}_h(\sigma_i, \sigma_{-i}) = \rho\}$. Here, $\mathbf{Out}_h(S_i)$ is the set of outcomes that are *compatible with* S_i. All notations for outcomes are lifted to histories in the obvious way. For a strategy profile $\sigma_P \in \Sigma$, we write $\mathbf{Hist}_h(\sigma_P)$ for the set $\{\rho_{\le j} \mid \rho \in \mathbf{Out}_h(\sigma_P), j \ge |h|\}$.

Consider two strategies σ and τ for player i, and a history h. We denote by $\sigma\left[h \leftarrow \tau\right]$ the strategy that follows strategy σ and *shifts* to τ at history h.

Formally, given a history h' such that $\mathrm{last}(h') \in V_i$:

$$\sigma\left[h \leftarrow \tau\right](h') \stackrel{\mathsf{def}}{=} \begin{cases} \tau(h^{-1} \cdot h') & \text{if } h \subseteq_{\mathsf{pref}} h' \\ \sigma(h') & \text{otherwise;} \end{cases}$$

We now formally define dominance and admissibility. We recall the intuition: a player's strategy σ is dominated by another strategy σ' of his if σ' yields a payoff which is as good as that of σ against all strategies for the other players, and is strictly better against some of them. A strategy is admissible if no other strategy dominates it. More formally, we have:

Dominance. A strategy $\sigma_i \in \Sigma_i$ *very weakly dominates* strategy $\sigma'_i \in \Sigma_i$, written $\sigma_i \succcurlyeq \sigma'_i$, if $\forall \sigma_{-i} \in \Sigma_{-i}, \mathsf{payoff}_i\big(\mathbf{Out}_{v_{\mathsf{init}}}(\sigma'_i, \sigma_{-i})\big) \le \mathsf{payoff}_i\big(\mathbf{Out}_{v_{\mathsf{init}}}(\sigma_i, \sigma_{-i})\big)$. Strategy σ_i *weakly dominates* strategy σ'_i, written $\sigma \succ \sigma'$, if $\sigma \succcurlyeq \sigma'$ and $\neg(\sigma' \succcurlyeq \sigma)$. A strategy $\sigma \in \Sigma_i$ is weakly dominated if there exists $\sigma' \in \Sigma_i$ such that $\sigma' \succ \sigma$. A strategy that is not weakly dominated is *admissible*. We denote by $\mathfrak{A}_i(\mathcal{G})$ the set of all admissible strategies for player i in $\mathcal{G}$.

Our characterizations and algorithms are based on the notions of *cooperative* and *antagonistic values* of a history. The antagonistic value, denoted $\mathbf{aVal}_i(\mathcal{G}, h)$, is the maximum payoff that player i can secure from h in the worst case, *i.e.* against all strategies of other players. The cooperative value, denoted $\mathbf{cVal}_i(\mathcal{G}, h)$, is the best value player i can achieve from h with the help of other players. We also define a third type of value: the *antagonistic-cooperative value*, denoted $\mathbf{acVal}_i(\mathcal{G}, h)$, which is the maximum value player i can achieve in $\mathcal{G}$ with the help of other players while guaranteeing the antagonistic value of the current history h. Formal definitions follow.

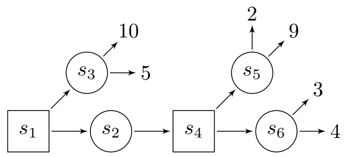

Figure 2 Example game where local conditions fail to capture admissibility.

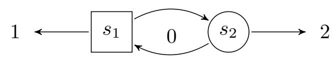

Figure 3 Example game with an infinite dominance chain and no admissible strategy as witness of their being dominated.

Antagonistic & Cooperative Values. The *antagonistic value of a strategy* and the *cooperative value of a strategy* σ_i of player i in $\mathcal{G}$, for a history h are

$$\mathbf{aVal}_i(\mathcal{G}, h, \sigma_i) \stackrel{\text{def}}{=} \inf_{\tau \in \Sigma_{-i}} \mathsf{payoff}_i\big(\mathbf{Out}_h(\sigma_i, \tau)\big);$$

$$\mathbf{cVal}_i(\mathcal{G}, h, \sigma_i) \stackrel{\text{def}}{=} \sup_{\tau \in \Sigma_{-i}} \mathsf{payoff}_i\big(\mathbf{Out}_h(\sigma_i, \tau)\big).$$

The *antagonistic value of a history* h for player i, and the *cooperative value of a history* h for player i are defined as $\mathbf{aVal}_i(\mathcal{G}, h) \stackrel{\text{def}}{=} \sup_{\sigma_i \in \Sigma_i} \mathbf{aVal}_i(\mathcal{G}, h, \sigma_i)$, and $\mathbf{cVal}_i(\mathcal{G}, h) \stackrel{\text{def}}{=} \sup_{\sigma_i \in \Sigma_i} \mathbf{cVal}_i(\mathcal{G}, h, \sigma_i)$, respectively. Finally, the *antagonistic-cooperative value of a history* h for player i is

$$\mathbf{acVal}_i(\mathcal{G}, h) \stackrel{\text{def}}{=} \sup\{\mathbf{cVal}_i(\mathcal{G}, h, \sigma_i) \mid \sigma_i \in \Sigma_i, \mathbf{aVal}_i(\mathcal{G}, h, \sigma_i) \geq \mathbf{aVal}_i(\mathcal{G}, h)\}.$$

We omit $\mathcal{G}$ when it is clear from the context.

Observe that $\mathbf{aVal}_i(h)$ of a history is the value of a zero-sum two-player game where player i is playing against players $-i$; while $\mathbf{cVal}_i(h)$ is the value in a one-player game, when all players play together. $\mathbf{acVal}_i(h)$ is a new notion which is the supremum of the values player i can obtain when he plays *worst-case optimal strategies.* A strategy $\sigma_i \in \Sigma_i$ is said to be worst-case optimal for player i at history h if $\mathbf{aVal}_i(h, \sigma_i) = \mathbf{aVal}_i(h)$; it is said to be *cooperatively optimal* for him at history h if $\mathbf{cVal}_i(h, \sigma_i) = \mathbf{cVal}_i(h)$. Observe that $\mathbf{acVal}_i(h) = -\infty$ if there are no worst-case optimal strategies from h.

▶ **Example 1** (Local conditions are not sufficient)**.** The game in Fig. 2 shows that admissibility requires one to consider the values of the histories both in the past and in the future of the current history. This shows that a local condition cannot capture admissibility. In fact, consider strategy σ_1 of player 1 (who controls all square vertices) that takes the edges $(s_1, s_2), (s_4, s_6)$. If the game starts at s_2, σ_1 is admissible, since the choice (s_4, s_5) could yield a payoff of 2 which is worse than any payoff from s_6. Indeed, we have that $\mathbf{aVal}_1(s_5) < \mathbf{aVal}_1(s_6)$. However, when the game starts at s_1, σ_1 is weakly dominated by the strategy that chooses (s_1, s_3) since the worst payoff in the latter case is 5. In fact, when a strategy takes the edge (s_1, s_2), the antagonistic value decreases from $\mathbf{aVal}_1(s_1) = 5$ to $\mathbf{aVal}_1(s_2) = 3$; so to be admissible, it should have a better cooperative value than 5, which is not the case if (s_4, s_6) is taken. The strategy taking $(s_1, s_2), (s_4, s_5)$ is admissible. Indeed, in one outcome, the payoff is 9, which is greater than 5 as required. Thus, an admissible strategy from s_1 either goes to s_3, or goes to s_2 but commits to taking (s_4, s_5) later.

We use temporal logic to describe sets of outcomes. We consider an extension of standard LTL with inequality conditions on payoffs for each player as in [3]. The logic, denoted $\mathsf{LTL}_{\mathsf{payoff}}$, extends LTL, and its syntax is defined as follows.

$$\varphi ::= Q \mid \neg\varphi \mid \mathtt{X}\varphi \mid \mathtt{G}\varphi \mid \mathtt{F}\varphi \mid \varphi_1 \,\mathtt{U}\, \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \mid \mathsf{payoff}_i \bowtie v,$$

Figure 4 Two games in which Player 1 has no admissible strategy.

where $Q \in \mathsf{AP}$ is a set of atomic propositions on edges, $\mathtt{G}$ and $\mathtt{F}$ are the standard LTL modalities, $\bowtie \in \{\leq, \geq, <, >\}$, and $v \in \mathbb{Q}$. A formula is interpreted over an outcome ρ at index k as follows. We have, for instance, $(\rho, k) \models Q$ if, and only if, (ρ_k, ρ_{k+1}) is labelled with Q. For convenience, we write $\rho \models \varphi$ instead of $(\rho, 1) \models \varphi$. Note that we define our predicates on edges rather than vertices; this simplifies our presentation. The semantics of the LTL modalities are standard; we refer to *e.g.* [3]. For payoff conditions, we have $(\rho, k) \models \mathsf{payoff}_i \bowtie v \overset{\text{def}}{\iff} \mathsf{payoff}_i(\rho_{\geq k}) \bowtie v$.

Residual Games. Given game $\mathcal{G}$, and history h, let us define $\mathcal{G}_h$ as the *residual game of $\mathcal{G}$ from h* by modifying the initial state to $\text{last}(h)$, and the payoff functions to payoff_i' defined as follows. For all outcomes ρ that start in $\text{last}(h)$, $\mathsf{payoff}_i'(\rho) = \mathsf{payoff}_i(h ^\frown \rho)$, where $h ^\frown \rho = h_{\leq |h|-1} \cdot \rho$. Notice that the strategy sets of $\mathcal{G}$ and $\mathcal{G}_h$ are identical, and that for any $\sigma_P \in \Sigma_P$, we have $\mathbf{Out}(\mathcal{G}_h, \sigma_P) = \mathbf{Out}_{\text{last}(h)}(\mathcal{G}, \sigma_P)$.

▶ **Lemma 2.** *For all $h' \in \mathbf{Hist}_{\text{last}(h)}(\mathcal{G})$, it holds that $\mathbf{aVal}_i(\mathcal{G}_h, h') = \mathbf{aVal}_i(\mathcal{G}, h ^\frown h')$, $\mathbf{acVal}_i(\mathcal{G}_h, h') = \mathbf{acVal}_i(\mathcal{G}, h ^\frown h')$, and $\mathbf{cVal}_i(\mathcal{G}_h, h') = \mathbf{cVal}_i(\mathcal{G}, h ^\frown h')$.*

3 Existence of Admissible Strategies

We start this section with two examples of quantitative games with no admissible strategies (for player 1). Then we identify a large and natural class of games for which the existence of admissible strategies is guaranteed.

Consider the games $\mathcal{A}$ and $\mathcal{G}$ in Fig. 4. Starting at s_1, the payoff of player 1, in the two games is defined as follows: an outcome that does not visit s_3 has a payoff equal to 0, otherwise, the payoff is equal to the number of times vertex a appears in the outcome. The lemma below states that player 1 does not have admissible strategies in those two games. We sketch the proof idea.

Consider first the one-player game $\mathcal{A}$. The antagonistic value at vertex s_1 is ∞. Any strategy which never visits s_3 is weakly dominated by strategies that visit a at least once (i.e. with outcome $(s_1 a s_1)^+ s_3^\omega$). Furthermore, a strategy which does visit s_3 and k times a is weakly dominated by any strategy that visits a at least $k+1$ times and then goes to s_3.

The idea is similar for $\mathcal{G}$ where the cooperative value at s_1 is ∞. Every strategy which does not allow outcomes visiting s_3 are weakly dominated by those that attempt to visit a by visiting s_2 at least once (as from s_2, the other player can cooperate and visit a), and then go to s_3. Moreover, it is always possible to attempt to visit a once more before going to s_3, thus any strategy which eventually goes to s_3 is also weakly dominated.

▶ **Lemma 3.** *Player 1 does not have admissible strategies in games $\mathcal{G}$ and $\mathcal{A}$.*

In the two examples above, either the $\mathbf{aVal}$ or the $\mathbf{cVal}$ (which are both equal to ∞) are not achievable. This is not a coincidence. We now show that all the games that admit witnessing strategies for those values are guaranteed to have admissible strategies.

Games with strategies witnessing aVal and cVal. A game is *well-formed* whenever it admits witnessing strategies for $\mathbf{aVal}$ and $\mathbf{cVal}$, *i.e.* it satisfies:

1. For all $i \in P$, and $h \in \mathbf{Hist}_{v_{\mathsf{init}}}(\mathcal{G})$, $\exists \sigma_i \in \Sigma_i, \mathbf{aVal}_i(h, \sigma_i) = \mathbf{aVal}_i(h)$.
2. For all $i \in P$, and $h \in \mathbf{Hist}_{v_{\mathsf{init}}}(\mathcal{G})$, $\exists \sigma_i \in \Sigma_i, \mathbf{cVal}_i(h, \sigma_i) = \mathbf{cVal}_i(h)$.

These conditions will also be referred as Assumption 1 and 2.

We now establish the existence of admissible strategies for all well-formed games.

▶ **Theorem 4.** *In all well-formed games all players have admissible strategies.*

The result follows from Lemmas. 6 and 7 below: the proof consists in showing that a particular type of admissible strategies, called *strongly cooperative-optimal*, always exists. Usually, those strategies are only a strict subset of the admissible strategies available to a player. Nevertheless, they are peculiar as they are guaranteed to exist.

▶ **Definition 5.** A strategy σ_i is *strongly cooperative-optimal (SCO)* if for all $h \in \mathbf{Hist}_{v_{\mathsf{init}}}(\sigma_i)$, if $\mathbf{cVal}_i(h) > \mathbf{aVal}_i(h)$ then $\mathbf{cVal}_i(h, \sigma_i) = \mathbf{cVal}_i(h)$, and if $\mathbf{aVal}_i(h) = \mathbf{cVal}_i(h)$ then $\mathbf{aVal}_i(h, \sigma_i) = \mathbf{aVal}_i(h)$.

Strongly cooperative-optimal strategies are admissible because their cooperative values are always maximal, and moreover, if a payoff better than the antagonistic value cannot be achieved ($\mathbf{aVal}_i(h) = \mathbf{cVal}_i(h)$), then they are worst-case optimal. Any strategy which obtains a better payoff than a SCO strategy against some adversary will obtain a worse payoff against another one.

▶ **Lemma 6.** *All strongly cooperative-optimal strategies are admissible.*

Proof. Let σ_i be a strongly cooperative-optimal strategy for player i. Assume towards a contradiction that some σ_i' weakly dominates σ_i. Let h be any minimal history compatible with σ_i such that $\sigma_i(h) \neq \sigma_i'(h)$.

If $\mathbf{aVal}_i(h) < \mathbf{cVal}_i(h)$, then since last$(h)$ is controlled by player i, $\mathbf{aVal}_i(h\sigma_i'(h)) \leq \mathbf{aVal}_i(h) < \mathbf{cVal}_i(h)$, and since σ_i is strongly cooperative optimal $\mathbf{cVal}_i(h\sigma_i(h), \sigma_i) = \mathbf{cVal}_i(h)$. Therefore, as the histories $h\sigma_i(h)$ and $h\sigma_i'(h)$ are distinct, there is a strategy $\tau \in \Sigma_{-i}$ such that $\mathsf{payoff}_i(\mathbf{Out}_{h\sigma_i(h)}(\sigma_i, \tau)) = \mathbf{cVal}_i(h) > \mathbf{aVal}_i(h) \geq \mathsf{payoff}_i(\mathbf{Out}_{h\sigma_i'(h)}(\sigma_i', \tau))$. This contradicts that σ_i' weakly dominates σ_i.

Otherwise $\mathbf{aVal}_i(h) = \mathbf{cVal}_i(h)$, then since σ_i is strongly cooperative optimal, for all $\tau \in \Sigma_{-i}$, $\mathsf{payoff}_i(\mathbf{Out}_h(\sigma_i, \tau)) = \mathbf{cVal}_i(h)$ and $\mathsf{payoff}_i(\mathbf{Out}_h(\sigma_i', \tau)) \leq \mathbf{cVal}_i(h)$. It follows that no outcome of σ_i' obtains a better payoff than σ_i. We thus obtain a contradiction. ◀

By Lem. 6, to prove the existence of admissible strategies, it suffices to prove the existence of strongly cooperative-optimal strategies. We actually give a constructive proof.

▶ **Lemma 7.** *In all well-formed games all players have SCO strategies.*

Let us describe the idea of the construction. Consider any player i. We define the strategy σ of player i as follows. For any history h, if $\mathbf{aVal}_i(h) = \mathbf{cVal}_i(h)$, then σ plays a worst-case optimal strategy from h, say σ_h^{wco}. Otherwise, we define σ starting from an outcome ρ_h with $\mathsf{payoff}_i = \mathbf{cVal}_i(h)$, and we define σ is such a way that it follows ρ_h. In this case, whenever another player deviates from ρ_h, say, at history h', we reevaluate how to play according to whether $\mathbf{aVal}_i(h') < \mathbf{cVal}_i(h')$ or $\mathbf{aVal}_i(h') = \mathbf{cVal}_i(h')$. Here, the existence of σ_h^{wco} and that of ρ_h are guaranteed by the fact that the game is well-formed.

In subsequent sections, we consider SCO strategies in residual games $\mathcal{G}_h$, so let us note that these games satisfy the required assumptions if $\mathcal{G}$ does, which follows from Lem. 2.

▶ **Lemma 8.** *For any well-formed game $\mathcal{G}$, for all histories $h \in \mathbf{Hist}_{v_{\mathsf{init}}}$, the residual game $\mathcal{G}_h$ is also well-formed.*

We end this section with an interesting observation: an infinite weak dominance chain is not necessarily dominated by an admissible strategy, as shown in the next example. The reader should contrast the example with the fact that in the Boolean case all dominated strategies are dominated by an admissible strategy [2, Thm. 11].

▶ **Example 9** (Non-dominated weak dominance chains)**.** There are quantitative games that have infinite dominance chains and no "maximal" admissible strategy weakly dominating them. Consider the game depicted in Fig. 3. Denote by σ^k the strategy of player 1 (controlling square vertices) which consists in moving from s_1 to s_2 exactly k times, and then going left (unless payoff of 2 was reached in the meantime). Then for all $k \in \mathbb{N}$, σ^k is weakly dominated by σ^{k+1} because if the adversary decides to move right from s_2 at the $(k+1)$-th step, σ^{k+1} performs better than σ^k, and otherwise they yield identical outcomes. It follows that all strategies σ^k for $k \geq 0$, are dominated. Here, the only admissible strategy σ^∞ consists in looping in the cycle forever, which does not dominate any σ^k since if the adversary always moves left from s_2, then σ^∞ yields less than σ^k.

▶ Remark. Above, we have defined strongly cooperative-optimal strategies that favour cooperation whenever it can have an added value. We have established that those strategies are always admissible. There are other classes of strategies that are always admissible, and we define another interesting class here. A strategy σ_i is a *worst-case cooperative optimal strategy*, if for all $h \in \mathbf{Hist}_{v_{\mathsf{init}}}(\sigma_i)$: $\mathbf{aVal}_i(h, \sigma_i) = \mathbf{aVal}_i(h)$, and $\mathbf{cVal}_i(h, \sigma_i) = \mathbf{acVal}_i(h)$.

So those strategies ensure the worst-case value at all times and leave open the best cooperation possible under that worst-case guarantee.

▶ **Lemma 10.** *All worst-case cooperative optimal strategies strategies are admissible.*

However, some well-formed games do not have worst-case cooperative optimal strategies.

4 Value-based Characterization of Admissible Strategies

We present our main result, which is, a value-based characterization of admissible strategies.

For any game $\mathcal{G}$, and player i, let us define the following property, denoted $\star(h, \sigma)$, for a given strategy $\sigma \in \Sigma_i(\mathcal{G})$ and history h:

$$\mathbf{cVal}_i(h, \sigma) > \mathbf{aVal}_i(h) \tag{1}$$
$$\vee \quad \mathbf{aVal}_i(h, \sigma) = \mathbf{cVal}_i(h, \sigma) = \mathbf{aVal}_i(h) = \mathbf{acVal}_i(h), \tag{2}$$

Intuitively, we will show that a strategy is admissible if at all histories, either the strategy promises a cooperative value greater than the antagonistic value at the current vertex, or a higher cooperative value cannot be obtained without risking a lower antagonistic value (*i.e.* $\mathbf{aVal}_i(h) = \mathbf{acVal}_i(h)$) and the strategy is worst-case optimal.

It turns out that requiring this property at all histories ending in a player's vertices characterize admissible strategies. We state our result in the following theorem.

▶ **Theorem 11.** *Under Assumption 1, for any game $\mathcal{G}$, player i, and $\sigma_i \in \Sigma_i(\mathcal{G})$, σ_i is admissible if, and only if, for all $h \in \mathbf{Hist}_{v_{\mathsf{init}}}(\mathcal{G}, \sigma_i)$ with $\mathrm{last}(h) \in V_i$, $\star(h, \sigma_i)$ holds.*

It will be useful to consider the negation of $\star(h, \sigma)$, which we simplify as follows:

▶ **Lemma 12.** *For all histories h and strategy σ, the negation of $\star(h,\sigma)$ is equivalent to*

$$\mathbf{cVal}_i(h,\sigma) \le \mathbf{aVal}_i(h) \wedge \mathbf{aVal}_i(h,\sigma) < \mathbf{aVal}_i(h) \quad (3)$$
$$\vee \quad \mathbf{cVal}_i(h,\sigma) = \mathbf{aVal}_i(h,\sigma) = \mathbf{aVal}_i(h) \wedge \mathbf{acVal}_i(h) > \mathbf{aVal}_i(h). \quad (4)$$

Proof of Thm. 11. $\boxed{\Rightarrow}$ We prove the contrapositive. Assume that $\exists h \in \mathbf{Hist}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i)$, $\mathsf{last}(h) \in V_i$ and $\neg\star(h,\sigma_i)$. Then by Lem. 12, either (3) or (4) holds for (h,σ_i).

Assume (3) holds for (h,σ_i). By Assumption 1, there exists a *worst-case optimal* strategy σ_h^{wco} from h, with $\mathbf{aVal}_i(h,\sigma_h^{\mathsf{wco}}) = \mathbf{aVal}_i(h)$. Define $\sigma_i' \stackrel{\mathsf{def}}{=} \sigma_i[h \leftarrow \sigma_h^{\mathsf{wco}}]$. We claim that σ_i' weakly dominates σ_i. In fact, for any $\sigma_{-i} \in \Sigma_{-i}(\mathcal{G})$ with $h \notin \mathbf{Hist}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_{-i})$, we have $\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i,\sigma_{-i}) = \mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i',\sigma_{-i})$. For any $\sigma_{-i} \in \Sigma_{-i}(\mathcal{G})$ compatible with h, both outcomes go through h. By definition of σ_i', $\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i',\sigma_{-i}) = h_{\le|h|-1} \cdot \mathbf{Out}_h(\mathcal{G},\sigma_h^{\mathsf{wco}},\sigma_{-i})$. Therefore, we have that $\mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i',\sigma_{-i})) \ge \mathbf{aVal}_i(h)$ by definition of σ_h^{wco}. The latter is greater than $\mathbf{cVal}_i(h,\sigma_i)$ from (3), so greater than $\mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i,\sigma_{-i}))$ by definition of $\mathbf{cVal}_i(\cdot)$. Thus, σ_i' very weakly dominates σ_i. Since by assumption, $\mathbf{aVal}_i(h,\sigma_i) < \mathbf{aVal}_i(h)$, and h is compatible with σ_i, there is a strategy $\sigma_{-i} \in \Sigma_{-i}$ such that $h \subseteq_{\mathsf{pref}} \mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i,\sigma_{-i})$ and $\mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i,\sigma_{-i})) < \mathbf{aVal}_i(h)$. As shown before, $\mathbf{aVal}_i(h) \le \mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i',\sigma_{-i}))$. Hence, σ_i' weakly dominates σ_i.

Assume now that (4) holds. Consider $\varepsilon > 0$ small enough so that $\mathbf{acVal}_i(h) > \mathbf{aVal}_i(h)+\varepsilon$. By definition of $\mathbf{acVal}_i(h)$, there exists a strategy $\tau_i \in \Sigma_i$ such that $\mathbf{cVal}_i(h,\tau_i) \ge \mathbf{aVal}_i(h)+\varepsilon$, and moreover $\mathbf{aVal}_i(h,\tau_i) \ge \mathbf{aVal}_i(h)$. Consider $\tau_{-i} \in \Sigma_{-i}$ compatible with h such that $\mathsf{payoff}_i(\mathbf{Out}_h(\mathcal{G},h,(\tau_i,\tau_{-i}))) \ge \mathbf{cVal}_i(h,\tau_i) - \frac{\varepsilon}{2} \ge \mathbf{aVal}_i(h) + \frac{\varepsilon}{2} > \mathbf{aVal}_i(h)$. Note that such a τ_{-i} exists by definition of $\mathbf{cVal}_i(h,\tau_i)$. It follows that $\sigma_i[h \leftarrow \tau_i]$ weakly dominates σ_i. In fact, the outcomes are identical for any outcome not compatible with h. For any σ_{-i} compatible with h, we have $\mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i,\sigma_{-i})) = \mathbf{aVal}_i(h)$ by (4). Moreover, $\mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i',\sigma_{-i})) \ge \mathbf{aVal}_i(h)$ since at h we have that $\mathbf{aVal}_i(h,\tau_i) \ge \mathbf{aVal}_i(h)$; thus $\mathbf{aVal}_i(h,\sigma_i') \ge \mathbf{aVal}_i(h)$. Furthermore, we have $\mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i',\tau_{-i})) > \mathbf{aVal}_i(h) \ge \mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i,\tau_{-i}))$.

$\boxed{\Leftarrow}$ Assume that for all $h \in \mathbf{Hist}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i)$ with $\mathsf{last}(h) \in V_i$, we have $\star(h,\sigma_i)$, and that σ_i is weakly dominated by some strategy σ_i'. We will show a contradiction.

Let σ_{-i} be a strategy in $\Sigma_{-i}(\mathcal{G})$ and $\rho = \mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i,\sigma_{-i})$ and $\rho' = \mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i',\sigma_{-i})$. If $\rho = \rho'$ then $\mathsf{payoff}_i(\rho') \le \mathsf{payoff}_i(\rho)$ and otherwise let j be the first index where they differ, and $h = \rho_{\le j-1} = \rho'_{\le j-1}$. We have that h is compatible with both strategies, $\mathsf{last}(h) \in V_i$ and $\sigma_i(h) \ne \sigma_i'(h)$.

If (1) holds, that is, $\mathbf{cVal}_i(h,\sigma_i) > \mathbf{aVal}_i(h)$, consider $\varepsilon > 0$ such that $\mathbf{cVal}_i(h,\sigma_i) > \mathbf{aVal}_i(h) + \varepsilon$, and a strategy $\sigma'_{-i} \in \Sigma_{-i}$ which ensures that $\mathsf{payoff}_i(\mathbf{Out}_{h\sigma_i(h)}(\mathcal{G},\sigma_i,\sigma'_{-i})) \ge \mathbf{cVal}_i(h,\sigma_i) - \frac{\varepsilon}{2}$, and $\mathsf{payoff}_i(\mathbf{Out}_{h\sigma_i'(h)}(\mathcal{G},\sigma_i',\sigma'_{-i})) \le \mathbf{aVal}_i(h,\sigma_i') + \frac{\varepsilon}{2}$. Such a strategy profile σ'_{-i} exists since $h\sigma_i'(h)$ and $h\sigma_i(h)$ are distinct, and since $\mathsf{last}(h) \in V_i$. The latter also implies that $\mathbf{aVal}_i(h) \ge \mathbf{aVal}_i(h,\sigma_i')$. It thus follows that

$$\mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i,\sigma_{-i}[h \leftarrow \sigma'_{-i}])) > \mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i',\sigma_{-i}[h \leftarrow \sigma'_{-i}]))$$

contradicting the fact that σ_i' weakly dominates σ_i.

Therefore (2) must hold, and $\mathbf{acVal}_i(h) = \mathbf{aVal}_i(h)$. If there exists $j \ge |h|$ such that $\mathbf{aVal}_i(\rho'_{\le j}) < \mathbf{aVal}_i(h)$, then there exists $\varepsilon > 0$ and a strategy profile $\sigma'_{-i} \in \Sigma_{-i}$ compatible with h which ensures that $\mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i',\sigma'_{-i})) \le \mathbf{aVal}_i(\rho'_{\le j}) + \varepsilon < \mathbf{aVal}_i(h) \le \mathsf{payoff}_i(\mathbf{Out}_{v_{\mathsf{init}}}(\mathcal{G},\sigma_i,\sigma'_{-i}))$. This contradicts σ' weakly dominating σ. Hence for all $j \ge |h|$, $\mathbf{aVal}_i(\rho'_{\le j}) \ge \mathbf{aVal}_i(h)$. Now, observe that $\mathsf{payoff}_i(\rho') \le \mathbf{acVal}_i(h)$. In fact, one can construct a strategy τ, which, from h follows ρ', and in case another player does not respect ρ, switches to a worst-case optimal strategy ensuring $\mathbf{aVal}_i(\rho'_{\le j}) \ge \mathbf{aVal}_i(h)$. It follows that $\mathsf{payoff}_i(\rho') \leq \mathbf{cVal}_i(h, \tau) \leq \mathbf{acVal}_i(h)$. Furthermore, by (2), $\mathsf{payoff}_i(\rho) \geq \mathbf{acVal}_i(h) = \mathbf{aVal}_i(h, \sigma_i)$, so $\mathsf{payoff}_i(\rho') \leq \mathsf{payoff}_i(\rho)$. This being true for all strategies of Σ_{-i} proves that σ_i very weakly dominates σ'_i and contradicts that σ'_i weakly dominates σ_i. ◀

5 Characterization of the Outcomes of Admissible Strategies

Observe that the characterization of Thm. 11 does not immediately yield an effective representation of the *set* of admissible strategies. In order to reason about the possible behaviors observable in a game under admissible strategies we are interested in describing the set of outcomes that can be observed when all players play admissible strategies. In this section, for each player, we give a linear temporal logic description of the outcomes that are each compatible with at least one admissible strategy.

Note that our main goal is to obtain such a characterization in full generality, for all well-formed games so we defer computability considerations to the next section. We will then see how the three types of values can be computed at all histories.

Let us fix a game $\mathcal{G}$, and player i. We present the intuition of the characterization. If an outcome ρ is compatible with an admissible strategy, say σ_i, then all prefixes h with $\mathrm{last}(h) \in V_i$ must satisfy (1) or (2). Given h, if (1) holds, then two things can happen. Either $\mathsf{payoff}_i(\rho) > \mathbf{aVal}_i(h)$, and thus ρ witnesses $\mathbf{cVal}_i(h, \sigma_i) > \mathbf{aVal}_i(h)$, or this is not the case but there is another outcome ρ' – compatible with σ_i – extending h with $\mathsf{payoff}_i(\rho') > \mathbf{aVal}_i(h)$. Notice how the longest common prefix of ρ and ρ' ends always with a vertex in V_{-i} since both outcomes are compatible with σ_i. If (2) holds at h, then, in particular, $\mathsf{payoff}_i(\rho) = \mathbf{aVal}_i(h)$ and, moreover, $\mathbf{aVal}_i$ remains constant at all prefixes of ρ extending h. The last observation simply follows from $\mathbf{aVal}_i(h, \sigma_i) = \mathbf{aVal}_i(h)$ which is implied by (2).

Extended $\mathsf{LTL}_{\mathsf{payoff}}$. Let $\mathbf{aValues}_i = \{\mathbf{aVal}_i(h) \mid h \text{ is a history}\}$ be the set of antagonistic values of player i. We will now define atomic propositions attached to edges of a game. Formally, we have a *labelling function* $\lambda : E \to \mathcal{P}(\mathsf{AP})$ which assigns to every edge a set of propositions from AP. The set AP includes the proposition $\mathsf{V_i}$ whose truth value, for every edge $e = (u, v)$, is determined as follows: $\mathsf{V_i} \in \lambda(u, v) \overset{\text{def}}{\iff} u \in V_i$.

We consider $\mathsf{LTL}_{\mathsf{payoff}}$ with atomic propositions as defined above and additional propositions aVal^i_q, acVal^i_q, and gAlt^i_q defined for all $q \in \mathbf{aValues}_i$. The semantics of these are straightforward: for an outcome ρ and $k \in \mathbb{N}_{>0}$ we have

$$(\rho, k) \models \mathsf{aVal}^i_q \overset{\text{def}}{\iff} \mathbf{aVal}_i(\rho_{\leq k}) = q,$$
$$(\rho, k) \models \mathsf{acVal}^i_q \overset{\text{def}}{\iff} \mathbf{acVal}_i(\rho_{\leq k}) = q, \text{ and}$$
$$(\rho, k) \models \mathsf{gAlt}^i_q \overset{\text{def}}{\iff} \rho_k \in V_{-i} \land \exists v' \neq \rho_k : (\rho_k, v') \in E \land \mathbf{cVal}_i(\rho_{\leq k} \cdot v') > q,$$

with the convention that, when k is omitted, we assume it is 1.

As mentioned earlier, we consider two cases depending on whether (1) or (2) hold. Thus, let us define the corresponding two sub-formulas:

$$\varphi_1 \overset{\text{def}}{=} \bigvee_{q \in \mathsf{aValues}_i} \left(\mathsf{aVal}^i_q \land \left(\mathsf{payoff}_i > q \lor \mathtt{F}(\mathsf{gAlt}^i_q)\right)\right), \text{ and}$$

$$\varphi_2 \overset{\text{def}}{=} \bigvee_{q \in \mathsf{aValues}_i} \left(\mathsf{aVal}^i_q \land \mathsf{acVal}^i_q \land \mathsf{payoff}_i = q \land \mathtt{G}\left(\mathsf{aVal}^i_q\right)\right).$$

We define the following formula which will be shown to capture the outcomes of admissible strategies: $\Phi^i_{\mathsf{adm}} \overset{\text{def}}{=} \mathtt{G}\left(\neg \mathsf{V_i} \lor \varphi_1 \lor \varphi_2\right)$.

▶ **Theorem 13.** *For any well-formed game $\mathcal{G}$, outcome ρ satisfies Φ^i_{adm} if, and only if, it is compatible with an admissible strategy for player i.*

We give the idea of the proof. For any outcome ρ compatible with an admissible strategy σ_i for player i. We show that for any prefix h of ρ with $\text{last}(h) \in V_i$, $(\rho, |h|)$ satisfies either φ_1 or φ_2. In fact, by Thm. 11, either (1) or (2) hold, and we show that these correspond to φ_1 and φ_2.

Conversely, for any ρ satisfying Φ^i_{adm}, we construct an admissible strategy σ_i for player i compatible with ρ. The strategy follows ρ, and in case of deviation, it switches immediately either to an SCO – which is guaranteed to exist – or to a worst-case optimal strategy, depending on whether φ_1 or φ_2 holds at the current history. The resulting strategy is proven to be admissible.

Assuming prefix-independence. Before concluding this section, let us consider the consequences of further assuming that our payoff function is *prefix-independent*.

3. For all $i \in P$, for all outcomes ρ, it holds that $\forall j \in \mathbb{N}, \mathsf{payoff}_i((\rho_k)_{k \geq j}) = \mathsf{payoff}_i(\rho)$.

Observe that, under Assumption 3, the set $\mathbf{aValues}_i$ can be equivalently defined as $\{\mathbf{aVal}_i(v) \mid v \in V\}$ and is thus finite. One can also extend the labelling λ and set of atomic propositions AP such that, for every edge $e = (u, v)$ and $q \in \mathbf{aValues}_i$:

$$\begin{aligned}
\mathsf{aVal}^i_q \in \lambda(u,v) &\overset{\text{def}}{\iff} \mathbf{aVal}_i(u) = q,\\
\mathsf{acVal}^i_q \in \lambda(u,v) &\overset{\text{def}}{\iff} \mathbf{acVal}_i(u) = q, \text{ and}\\
\mathsf{gAlt}^i_q \in \lambda(u,v) &\overset{\text{def}}{\iff} u \in V_{-i} \wedge \exists v' \neq v : (u,v') \in E \wedge \mathbf{cVal}_i(v') > q.
\end{aligned}$$

It immediately follows that:

▶ **Lemma 14.** *Under Assumption 3, for all $i \in P$, Φ^i_{adm} is expressible in $\mathsf{LTL}_{\mathsf{payoff}}$.*

6 Applications and Future Works

In this section, we show how to apply Theorem 11 (value-based characterization of admissible strategies) and Theorem 13 (characterization of the set of outcomes of admissible strategies) to solve relevant verification and synthesis problems.

Classical payoff functions. So far, we have assumed that games were equipped for each player $i \in P$ with a payoff function. To define payoff functions, we proceed as usual by first assigning weights to edges of the game graph using *weight functions* $w_i : E \to \mathbb{Q}$, one for each player $i \in P$. With the weight function w_i, we associate to each outcome $\rho = \rho_1\rho_2 \ldots \rho_n \ldots$, an infinite sequence of rational values $w_i(\rho) = w_i(\rho_1\rho_2)w_i(\rho_2\rho_3) \ldots w_i(\rho_n\rho_{n+1}) \ldots$, and we aggregate this sequence of values with measures such as Inf, Sup, LimInf, LimSup, and mean payoff ($\underline{\mathsf{MP}}$ and $\overline{\mathsf{MP}}$). It is well known, see *e.g.* [7] and [18], that all the payoff functions defined above satisfy Assumptions 1-2. By Theorem 4, we get the following.

▶ **Lemma 15.** *In games with payoff functions from Inf, Sup, LimInf, LimSup, $\underline{\mathsf{MP}}$, and $\overline{\mathsf{MP}}$, all players have admissible strategies.*

It is also known that, in games defined with the payoff functions considered here, the antagonistic and cooperative values ($\mathbf{cVal}$ and $\mathbf{aVal}$) are computable. One can also show that $\mathbf{acVal}$ is computable for prefix-independent payoff functions. Indeed, this value of a vertex coincides with the **cVal** inside the sub-graph induced by the vertices with the optimal antagonistic value. Furthermore, using a classical transformation on the game structure, we can guarantee that all payoff functions We thus obtain the following result, by Lemma 14.

▶ **Lemma 16.** *In games with payoff functions from* Inf*,* Sup*,* LimInf*,* LimSup*,* $\underline{\mathsf{MP}}$*, and* $\overline{\mathsf{MP}}$*, the formulas* Φ^i_{adm} *for all* $i \in P$ *are effectively computable, finite, and expressible in* $\mathsf{LTL}_{\mathsf{payoff}}$*.*

We will now consider several problems of interest which can be solved using the characterizations that we have developed in the previous sections. All the results are applicable to the measures concerned by Lemmas 15 and 16.

Deciding the admissibility of a finite memory strategy. As a first example, we consider the problem of deciding, given a game structure $\mathcal{G}$, and a (finite memory) strategy σ_i for player $i \in P$ described as a finite state transducer M_i, if σ_i is admissible in $\mathcal{G}$.

To solve this problem, we rely on Theorem 11 and proceed as follows. First, we compute for each vertex v of the game $\mathcal{G}$, the values $\mathbf{aVal}_i(\mathcal{G}, v)$, $\mathbf{cVal}_i(\mathcal{G}, v)$, and $\mathbf{acVal}_i(\mathcal{G}, v)$. Second, we construct the synchronized product between the transducer M_i that defines the strategy σ_i and the game $\mathcal{G}$. States in this product are of the form (v, m) where v is a vertex of $\mathcal{G}$ and m is a (memory) state of the transducer M_i. Third, we compute for each state (v, m) the values $\mathbf{aVal}_i(\mathcal{G}, (v, m), \sigma_i)$, $\mathbf{cVal}_i(\mathcal{G}, (v, m), \sigma_i)$, and $\mathbf{acVal}_i(\mathcal{G}, (v, m), \sigma_i)$. Finally, we verify that there is no reachable vertex (v, m) in the product where condition (1) or condition (2) are falsified. We then obtain the following theorem:

▶ **Theorem 17.** *Given a game* $\mathcal{G}$ *and a finite memory strategy* σ_i *for player* $i \in P$ *specified as a finite state transducer* M_i*, we can decide if* σ_i *is an admissible strategy for player* i *in* PTime *for measures* Inf*,* Sup*,* LimInf*,* LimSup*; in* NP ∩ coNP *for* $\underline{\mathsf{MP}}$*, and* $\overline{\mathsf{MP}}$*.*

Model-checking under admissibility. We now turn to the following problem. Given a game structure $\mathcal{G}$ and a $\mathsf{LTL}_{\mathsf{payoff}}$ formula φ, decide if all outcomes of the game that are compatible with the admissible strategies of all players satisfy φ, *i.e.* if $\bigcap_{i \in P} \mathbf{Out}(\mathcal{G}, \mathfrak{A}_i(\mathcal{G})) \models \varphi$. This problem was introduced in the Boolean setting in [5] and allows one to check that a property is induced by the rationality of the players in a game.

▶ **Theorem 18.** *For all measures* Inf*,* Sup*,* LimInf*,* LimSup*,* $\underline{\mathsf{MP}}$*,* $\overline{\mathsf{MP}}$*, one can decide, given game* $\mathcal{G}$ *and* $\mathsf{LTL}_{\mathsf{payoff}}$ *formula* φ*, whether* $\bigcap_{i \in P} \mathbf{Out}(\mathcal{G}, \mathfrak{A}_i(\mathcal{G})) \models \varphi$*.*

Proof Sketch. For each player $i \in P$, consider the formula Φ^i_{adm} from Theorem 13, which describes the set $\mathbf{Out}(\mathcal{G}, \mathfrak{A}_i(\mathcal{G}))$. The formula is finite and constructible by Lemma 16. The problem now amounts to verifying if $\mathcal{G}$ satisfies the specification $\left(\bigwedge_{i \in P} \Phi^i_{\mathsf{adm}}\right) \implies \varphi$. For all payoff functions, except mean-payoff, this can be reduced to model checking an LTL formula (since the measures are regular). For $\overline{\mathsf{MP}}$ and $\underline{\mathsf{MP}}$, the result follows from [3] which shows that the model checking problem against $\mathsf{LTL}_{\mathsf{payoff}}$ is decidable. ◀

Quantitative assume-admissible synthesis. In [4], a new rule for reactive synthesis in non-zero sum n-player games was proposed. The setting there is similar to the setting considered here but it is Boolean: each player $i \in P$ has his own omega-regular objective $O_i \subseteq V^\omega$. The synthesis rule asks if player $i \in P$ has a strategy to enforce its own objective O_i against admissible strategies of the other players. In other words, the rule asks for the existence of worst-case optimal strategies against rational adversaries.

The quantitative extension of this problem asks given a game $\mathcal{G}$, a player $i \in P$, and a $\mathsf{LTL}_{\mathsf{payoff}}$ formula φ, $\exists \sigma \in \mathfrak{A}_i, \forall \tau \in \mathfrak{A}_{-i}, \varphi$. Using Theorem 13, we can reduce this query to a plain two-player zero-sum game on the game structure $\mathcal{G}$ with objective:

$$\exists \sigma \in \Sigma_i, \forall \tau \in \Sigma_{-i}, \Phi^i_{\mathsf{adm}} \wedge \left(\left(\bigwedge_{j\in P\setminus\{i\}} \Phi^j_{\mathsf{adm}}\right) \implies \varphi\right)$$

Since for Inf, Sup, LimInf, LimSup, Φ^i_{adm} and φ are omega-regular, the problem reduces to deciding the winner in a two-player zero-sum game with omega-regular objectives. As a consequence, we obtain the following theorem:

▶ **Theorem 19.** *The quantitative assume-admissible synthesis problem for player $i \in P$ is decidable for measures* Inf*,* Sup*,* LimInf*,* LimSup*.*

For the measures $\underline{\mathsf{MP}}$, $\overline{\mathsf{MP}}$, we obtain objectives in which mean-payoff constraints and omega-regular constraints are mixed. On the one hand, those objectives are outside known decidable classes of objectives treated in [8] and in [10]. On the other hand, the undecidability results obtained in [17] do not apply to them. This motivates further research on zero-sum two player games with a mix of mean-payoff and omega-regular objectives.

Towards iterative elimination. Once we have computed the admissible strategies for each player, we restrict each player to these strategies, and repeat the computation of the admissible strategies in the restricted game. This can be iterated several times and gives a process that is called *iterative elimination of dominated strategies*, and well known in game theory. This process is difficult to analyze for mean-payoff, because objectives of different players interfere in non-trivial ways and games with Boolean combinations of mean-payoff objectives are undecidable [17]. However it seems feasible for regular payoffs, such as Inf, Sup, LimInf and LimSup, for which we can construct parity automata recognizing outcomes with $\mathsf{payoff}_i > q$. Given $i \geq 0$, we can actually compute a parity automaton accepting the set of outcomes of $\mathcal{S}^i$ which is the set of strategies that remain after i steps of elimination. We summarize here the ingredients but leave the details for future work. Assume we have a parity automaton representing the outcomes of $\mathcal{S}^i$. Note that for $i = 0$ this is simply all outcomes. If the payoffs are regular, then we can compute values $\mathbf{cVal}_i(h, \mathcal{S}^i)$, $\mathbf{aVal}_i(h, \mathcal{S}^i)$ and $\mathbf{acVal}_i(h, \mathcal{S}^i)$, which correspond to cooperative, antagonistic, and antagonist-cooperative values when players only play strategies from $\mathcal{S}^i$. We can then use these values as atomic propositions for a $\mathsf{LTL}_{\mathsf{payoff}}$ formulas similar to Φ^i_{adm} of Section 5, which characterizes outcomes of strategies of $\mathcal{S}^{i+1}$. In the case of regular payoffs this yields a parity automaton which represents the outcomes of $\mathcal{S}^{i+1}$. This procedure can then be repeated to compute outcomes that are possible under iterative elimination.

References

1 Brandenburger Adam, Friedenberg Amanda, H Jerome, et al. Admissibility in games. *Econometrica*, 2008.

2 Dietmar Berwanger. Admissibility in infinite games. In *STACS*, volume 4393 of *LNCS*, pages 188–199. Springer, February 2007.

3 Udi Boker, Krishnendu Chatterjee, Thomas A. Henzinger, and Orna Kupferman. Temporal specifications with accumulative values. *ACM Trans. Comput. Logic*, 15(4):27:1–27:25, July 2014. `doi:10.1145/2629686`.

4 Romain Brenguier, Jean-François Raskin, and Ocan Sankur. Assume-admissible synthesis. In Luca Aceto and David de Frutos-Escrig, editors, *CONCUR*, volume 42 of *LIPIcs*, pages

100–113. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2015. doi:10.4230/LIPIcs.CONCUR.2015.100.

5 Romain Brenguier, Jean-François Raskin, and Mathieu Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS'14, 2014*. ACM, 2014. doi:10.1145/2603088.2603143.

6 Krishnendu Chatterjee, Laurent Doyen, Emmanuel Filiot, and Jean-François Raskin. Doomsday equilibria for omega-regular games. In *VMCAI'14*, volume 8318, pages 78–97. Springer, 2014.

7 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Quantitative languages. *ACM Transactions on Computational Logic*, 11(4), 2010. doi:10.1145/1805950.1805953.

8 Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdziński. Mean-payoff parity games. In *LICS*, pages 178–187. IEEE Computer Society, 2005. doi:10.1109/LICS.2005.26.

9 Krishnendu Chatterjee, Thomas A Henzinger, and Marcin Jurdziński. Games with secure equilibria. *Theoretical Computer Science*, 365(1):67–82, 2006.

10 Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. Strategy synthesis for multi-dimensional quantitative objectives. *Acta Inf.*, 51(3-4):129–163, 2014. doi:10.1007/s00236-013-0182-6.

11 Marco Faella. Admissible strategies in infinite games over graphs. In *MFCS 2009*, volume 5734 of *Lecture Notes in Computer Science*, pages 307–318. Springer, 2009.

12 Dana Fisman, Orna Kupferman, and Yoad Lustig. Rational synthesis. In *TACAS'10*, volume 6015 of *LNCS*, pages 190–204. Springer, 2010.

13 R. J. Gretlein. Dominance elimination procedures on finite alternative games. *International Journal of Game Theory*, 12(2):107–113, 1983. doi:10.1007/BF01774300.

14 O. Kupferman, G. Perelli, and M.Y. Vardi. Synthesis with rational environments. In *Proc. 12th European Conference on Multi-Agent Systems*, LNCS. Springer, 2014.

15 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190, 1989. doi:10.1145/75277.75293.

16 Wolfgang Thomas. On the synthesis of strategies in infinite games. In *STACS*, pages 1–13, 1995. doi:10.1007/3-540-59042-0_57.

17 Yaron Velner. Robust multidimensional mean-payoff games are undecidable. In Andrew M. Pitts, editor, *FoSSaCS*, volume 9034 of *LNCS*, pages 312–327. Springer, 2015. doi:10.1007/978-3-662-46678-0_20.

18 Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *TCS*, 158(1):343–359, 1996.