

Admissibility in Infinite Games

Dietmar Berwanger*

RWTH Aachen, Mathematical Foundations of Computer Science
52056 Aachen, Germany, berwanger@cs.rwth-aachen.de

Abstract We analyse the notion of iterated admissibility, i.e., avoidance of weakly dominated strategies, as a solution concept for extensive games of infinite horizon. This concept is known to provide a valuable criterion for selecting among multiple equilibria and to yield sharp predictions in finite games. However, generalisations to the infinite are inherently problematic, due to unbounded dominance chains and the requirement of transfinite induction.

In a multi-player non-zero-sum setting, we show that for infinite extensive games of perfect information with only two possible payoffs (win or lose), the concept of iterated admissibility is sound and robust: all iteration stages are dominated by admissible strategies, the iteration is non-stagnating, and, under regular winning conditions, strategies that survive iterated elimination of dominated strategies form a regular set.

1 Introduction

Games are fundamental for the analysis of interaction between computational systems [19]. As a particularly effective model, sequential two-player zero-sum games of infinite duration are paradigmatic for capturing the nonterminating behaviour of reactive systems in interplay with their environment [22, 11]. For this class of games, a rich and powerful theory has been developed over the past fifty years: the semantics of most temporal and fixed-point logics translates into this framework, and the relevant algorithmic questions can be captured in terms of alternating ω -automata [20]. More recently, the research on multi-agent systems and on self-organising computational networks motivated the study of models with more than two players that are not necessarily in conflict. These investigations typically focus on the strategic choices of decision-makers rather than their sequential behaviour.

There are good reasons for investigating sequential games of infinite duration with more than two players. In supervisory control theory, the components of a distributed systems can be conceived as distinct players with the same payoff, rather than a coalition competing against the environment [2, 17]. In the context of autonomous agents, the framework captures models of infinite horizon. For the analysis of interaction in large networks, infinite multi-player games provide a basic model that takes into account the sequential, possibly nonterminating behaviour of the participating systems.

* This research was supported by the EU RTN 'GAMES' (www.games.rwth-aachen.de).

In this paper, we initiate a fundamental study addressing the question of rational behaviour in games of infinite duration that involve two or more players and are not strictly competitive. Concretely, we generalise the classical concept of iterated admissibility, that is well established in the theory of finite games [14, 7], to the framework of n -player non-zero-sum extensive games of perfect information over trees of countable depth. This concept is based on the notion of (weak) dominance. Considering two strategies r, s of a particular player, s dominates r if, against any choice of strategies of the other players, s performs at least as well as r , but there are cases in which s performs strictly better than r . As a basic principle of rationality it is assumed that, when a player takes all strategies of the other players in consideration, he will avoid playing dominated strategies. Accordingly, each player eliminates from his set of strategies those that are dominated. On the other hand, rationality is assumed to be common knowledge. Therefore each player understands which strategies his opponents eliminated, possibly discovering that some of his remaining strategies are again dominated when taking into account only admissible counter-strategies. This leads to the elimination of further strategies, and so on. For games with finite strategy spaces, the procedure stabilises after a finite number of stages; the solution concept postulating that outcomes of a game should involve only strategies that survive the iterated elimination is called iterated admissibility (a precise definition follows). Despite its simplicity, the analysis of this procedure required extensive efforts [9, 16, 1, 18]. An epistemic characterisation, justifying on the basis of the player's introspective abilities why outcomes should be iteratively admissible, was achieved recently [3], after having been open for a long time.

The generalisation of iterated admissibility to infinite strategy spaces is not obvious. Already at the elementary level, the avoidance of dominated strategies can lead to empty solutions, because maximally dominating strategies may not exist. For instance, in a game where two players simultaneously choose a natural number, with the greater number winning, any strategy choosing x is dominated by the strategy choosing $x + 1$. Unbounded dominance chains are problematic in general: it does not seem reasonable to eliminate a strategy unless some other strategy that dominates it survives. As a further difficulty, the elimination process may need infinitely many iterations to stabilise. In order to mitigate some of these apparent inconveniences, we will restrict our attention to games with qualitative payoffs, where a player may either win or lose. This setting covers a variety of situations significant in computer science.

We show that in the framework of infinite sequential games of perfect information with qualitative payoffs, the generalisation of iterated admissibility leads to a viable solution concept. For extensive games with winning conditions given by arbitrary sets of infinite paths, we prove that the procedure of simultaneous elimination of dominated strategies is sound, in the sense that every eliminated strategy is dominated by some surviving strategy. Further, we point out that the promoted solution is compatible with the sequential game structure: for any position reachable with admissible strategies, a strategy is admissible in the subgame rooted at that position if, and only if, it is the restriction of an admissible

strategy in the original game. In particular, this means that admissibility does not feature non-credible threats. Although the iteration may not terminate in the general case, we still show that the procedure does not stagnate, in the sense that it generates a definitive solution for a maximal antichain in the game tree in a bounded number of stages.

Finally, we discuss iterated admissibility as a solution concept for ω -regular non-zero-sum graph games with n players. These are games played on the tree unravelling of a finite labelled graph with winning conditions given by ω -regular languages over the alphabet consisting of these labels. This framework is relevant for the specification, verification, and synthesis of finite-state controllers. Due to the non-stagnation property, the iteration on regular graph games always terminates after a finite number of steps. Strategies can be conceived as labelled ω -trees of finite branching. We show that, if we set out with ω -tree-regular sets of strategies, one for each player, the set of undominated strategies for any player is again regular. As a consequence, it follows that the set of iteratively admissible strategies in an ω -regular graph game is regular.

The area of infinite sequential games with non-zero-sum conditions or more than two players is largely unexplored. Related research focusses on Nash equilibrium and its refinements [5, 6, 21, 4, 8]. Nash equilibrium is the most prominent solution concept in the classical theory of non-cooperative games. It postulates that the collective decision of the players should be self-enforcing, in the sense that no one gains if he alone deviates. This concept has proved to be highly effective in explaining the outcome of games involving economic agents. However, in predicting and prescribing rational behaviour, Nash equilibria suffer from several shortcomings, some of which aggravate in the context of computational agents. One fundamental problem concerns games with multiple equilibria, where the interplay of strategies that are each consistent with some equilibrium may yield non-equilibrium outcomes (coordination failure). Justifications that players would coordinate on a particular equilibrium point often resort to evolutionary arguments that do not apply to computational agents [15]. Another problem is that Nash equilibria disregard the sequential structure of extensive games (non-credible threats). Besides constituting a solution concept on its own right, iterated admissibility provides basic criteria to exclude such implausible Nash equilibria and to mitigate the effect of coordination failures [13, 12].

The paper is organised as follows. In Section 2 we introduce basic notions and define the solution concept of iterated admissibility. Our principal technical tool is a value characterisation of admissibility developed in Section 3. Given a profile of strategy sets, the value for a particular player is a colouring of game positions that reflects the expectations of the player – to surely win, possibly win, or surely lose – when the play reaches that position, assuming that all players use strategies from the given sets. Using this characterisation, we show in Section 4 that the elimination procedure is sound and, in Section 5, that it produces a solution for an antichain of subgames whenever it is iterated over a bounded number of times. Finally, Section 6 is dedicated to ω -regular graph games. There, we show that iterated admissibility preserves regularity.

2 Basic notions

We are concerned with games involving n players. We refer to a list of elements $x = (x^i)_{i < n}$, one for each player, as a *profile*. For any such profile, we write x^{-i} to denote the list $(x^j)_{j < n, i \neq j}$ of elements in x for all players except i . Given an element x^i and a list x^{-i} , we denote by (x^i, x^{-i}) the profile $(x^i)_{i < n}$. For clarity, we will always use superscripts to specify to which player an element belongs. If not quantified otherwise, we usually refer to player i meaning *any* player.

Definition 1. A game in *strategic form* is a structure $\Gamma = ((S^i)_{i < n}, (u^i)_{i < n})$ consisting of sets S^i of *strategies*, and *utility payoff* functions $u^i : \times_{i < n} S_i \rightarrow \mathbb{R}$, one for every player i .

When a game is played, each player i chooses a strategy $s^i \in S^i$, thus yielding a strategy profile $s = (s^i)_{i < n}$. The value $u^j(s)$ represents the payoff received by player j in this play. We denote by S the set of all strategy profiles $\times_{i < n} S^i$.

Every player is assumed to act rationally towards maximising his own payoff. A *solution concept* associates to any game a nonempty subset of strategy profiles that are consistent with this assumption. However, the notion of rationality is an informal one, allowing a variety of interpretations. We investigate a formalisation based on the notion of (weak) dominance.

Definition 2. Given two strategies $s^i, r^i \in S^i$ and a set $Q^{-i} \subseteq S^{-i}$, we say that s^i *dominates* r^i on Q^{-i} , if

$$\begin{aligned} u^i(s^i, t^{-i}) &\geq u^i(r^i, t^{-i}) \text{ for all } t^{-i} \in Q^{-i}, \text{ and} \\ u^i(s^i, t^{-i}) &> u^i(r^i, t^{-i}) \text{ for some } t^{-i} \in Q^{-i}. \end{aligned}$$

For a set $Q \subseteq S$, we say that a strategy $s^i \in S^i$ is *admissible with respect to* Q , if no strategy in Q^i dominates s^i on Q^{-i} .

The idea is to postulate first that it is rational to avoid dominated strategies. Assuming common knowledge of rationality, all players need to be aware of this, and iteratively reevaluate their strategy sets. Different non-equivalent procedures have been proposed on the basis of this idea. Following [3], we opt for simultaneous maximal elimination.

Definition 3. For a game Γ , we define simultaneously for all players i :

- $Q_0^i := S^i$;
- $Q_{\alpha+1}^i := \{s^i \in Q_\alpha^i : s^i \text{ is admissible w.r.t } Q_\alpha\}$, for every ordinal α , and
- $Q_\lambda^i := \bigcap_{\alpha < \lambda} Q_\alpha^i$ for every limit ordinal λ .

A strategy $s^i \in Q_\alpha^i$ is called α -*admissible*. As the stages are decreasing, the induction reaches a fixed point $(Q_\infty^i)_{i < n}$. A strategy $s^i \in Q_\infty^i$ is called *iteratively admissible*.

Observe that, in contrast to most equilibrium concepts, iterative admissibility yields a *rectangular* set of strategy profiles, i.e., a Cartesian product of sets. Accordingly, whether the choice of a particular player is rational in our sense does not depend on the choices of the other players.

Extensive games of infinite horizon. As illustrated in the introduction, there is no hope that admissibility yields a meaningful solution for arbitrary infinite games. We will henceforth restrict our attention to a particular class of sequential games with perfect information.

Definition 4. A game in *extensive form* is a structure $\Gamma = (\mathcal{T}, (W^i)_{i < n})$, where \mathcal{T} is a directed tree over a domain T of elements called *positions*, equipped with a partition $(T^i)_{i < n}$ of its non-terminal nodes, and each W^i is a set of maximal paths through \mathcal{T} , called the *winning set* of player i .

Game trees may be of arbitrary branching and of depth up to ω . We write \preceq for the partial order associated to \mathcal{T} . When the game is played, the players form a maximal path through \mathcal{T} starting from the root. Whenever a position p is reached, the player i to whose partition T^i it belongs, prolongs the path along some edge. Thus, *plays* are identified with maximal paths through \mathcal{T} . An *initial play* is a prefix of a play. The *length* of an initial play $\pi = p_0, p_1, \dots, p_\ell$ is ℓ .

Definition 5. A *strategy* for player i is a function $s : T^i \rightarrow T$ that associates to every position in T^i a successor in T .

Again, we denote the set of all strategies for player i by S^i . An (initial) play p_0, p_1, \dots follows a strategy $s^i \in S^i$, if for every $p_\ell \in T^i$, we have $p_{\ell+1} = s(p_\ell)$. A position $p \in T$ is *avoided* by s^i , if the unique initial play that ends at p does not follow s^i , otherwise p is *reachable* by s^i . Given a strategy set $Q^i \subseteq S^i$, we write $Q^i(p)$ for the subset of strategies in Q^i that do not avoid p . For $Q \subseteq S$, we define $Q(p)$ as the set of profiles $s \in Q$ with $s^i \in Q^i(p)$ for all $i < n$. We say that p is *reachable* in Q^i or in Q , if $Q^i(p)$ or $Q(p)$, respectively, is not empty. Two strategies $r^i, s^i \in S^i$ *agree* up to a position p , if for every $p' \preceq p$ with $p' \in T^i$, we have $r^i(p') = s^i(p')$; they *split* at p , if they agree up to p , but not on the immediate prolongation of the play, i.e., $r^i(p) \neq s^i(p)$.

Any strategy profile $s = (s^i)_{i < n}$ determines a unique play π which follows all of its components. This play is the *outcome* of the profile; in the 2-player case we denote it by $s^0 \hat{\ } s^1$. A profile is *winning* for player i if its outcome is in W^i . In terms of utility, we have $u^i(\pi) = 1$ iff $\pi \in W^i$, otherwise $u^i(\pi) = -1$.

Given a game Γ and a position p , the *subgame* Γ_p is the game obtained by restricting all components of Γ to the positions comparable to p . We denote the restriction of a strategy $s^i \in S^i$ to this domain by $s^i|_p$. Given a strategy set $Q^i \subseteq S^i$, we define $Q^i|_p := \{s^i|_p : s^i \in Q^i(p)\}$. For a profile $Q \subseteq S$ of strategy sets, the restriction is defined component-wise $Q|_p = (Q^i|_p)_{i < n}$.

3 Value characterisation

The main technical tool for our analysis is a characterisation of admissible strategies in terms of the value that a player can achieve in a given (restriction of the) strategy space. This characterisation also links rational behaviour in subgames with the original game. Whenever we refer to games in the sequel, we mean infinite extensive games.

Definition 6 (Value). Fix a game Γ and a rectangular set $Q \subseteq S$. The *value* of Q for player i is a function χ^i that maps every position p reachable in Q to an element of $\{-1, 0, 1\}$, as follows:

- $\chi^i(p) = 1$ (winning), if there exists a strategy $s^i \in Q^i(p)$ such that for any $t^{-i} \in Q^{-i}(p)$, the profile (s^i, t^{-i}) is winning for player i ;
- $\chi^i(p) = -1$ (losing), if every profile in $Q(p)$ is losing for player i ;
- $\chi^i(p) = 0$ (pending), otherwise.

First, we define a criterion that identifies well-formed restrictions of the strategy space in extensive games, which allows us to conceive a game position as a point where the player in turn may reassess his plan of action.

Definition 7 (Shifting). Given $s^i, r^i \in S^i$ and a position p up to which the two strategies agree, we denote by $s^i[p \leftarrow r^i]$ the strategy that agrees with r^i on every position comparable to p , and with s^i on any other position.

We say that a strategy set Q^i *allows shifting*, if for any $s^i, r^i \in Q^i$, and every position p up to which s^i and r^i agree, we have $s^i[p \leftarrow r^i] \in Q^i$. A rectangular set $Q \subseteq S$ of profiles allows shifting if all its components do.

The following lemma states that a strategy s dominates r if, and only if, from any position at which s and r disagree, either they both win or both lose, or s strictly improves over r .

Lemma 8. *Consider two strategies $s^i, r^i \in S^i$ and a rectangular set $Q \subseteq S$ that allows shifting. We denote the values of $\{s^i\} \times Q^{-i}$ and $\{r^i\} \times Q^{-i}$ for player i by χ_s^i and χ_r^i , respectively. Then, s^i dominates r^i on Q^{-i} if, and only if,*

- (i) *for any position p at which r^i and s^i split, we have $\chi_s^i(p) \geq \chi_r^i(p)$, but at most one of $\chi_s^i(p)$ and $\chi_r^i(p)$ is 0, and*
- (ii) *there exists a position p up to which s^i and r^i agree, where $\chi_s^i(p) > \chi_r^i(p)$.*

Proof. We only show here that the conditions are necessary, restricting to the case of two players with $i = 0$ and omitting the superscripts.

First, consider strategies $s, r \in S^0$ that violate condition (i), i.e., they split at some position p where the values are either $\chi_s(p) < \chi_r(p)$ or $\chi_s(p) = \chi_r(p) = 0$. In the former case, it follows that there exists a strategy $t \in Q^1$ that reaches p and further yields $s \hat{t} \notin W^0$ whereas $r \hat{t} \in W^0$. Hence s cannot dominate r on Q^1 . In the latter case, since $\chi_s(p) < 1$, there exists a strategy $t \in Q^1(p)$ such that $s \hat{t} \notin W^0$. On the other hand, since $\chi_r(p) > -1$, there also exists a strategy $t' \in Q^1(p)$ such that $r \hat{t}' \in W^0$. Let $p' := r(p)$ be the position chosen by r , differing from $s(p)$. Since t, t' agree up to p , and Q^1 allows shifting, the strategy $\tau := t[p' \leftarrow t']$ is also in Q^1 . But now $r \hat{\tau} = r \hat{t}' \in W^0$, whereas $s \hat{\tau} = s \hat{t} \notin W^0$, which shows that s does not dominate r on Q^1 . Whenever two strategies s, r satisfy the first condition but not the second one, they shall have equal, but nonzero values at any splitting position p : either $\chi_s(p) = \chi_r(p) = 1$, or $\chi_s(p) = \chi_r(p) = -1$. It follows that for every $t \in Q^1$ we have $s \hat{t} \in W^0$ if, and only if, $r \hat{t} \in W^0$. Hence the strategies s and r are incomparable on Q^1 . \square

The characterisation of dominance extends to a characterisation of admissible strategies.

Lemma 9. *Let $Q \subseteq S$ be a rectangular set that allows shifting. Then, a strategy $s^i \in S^i$ is admissible with respect to Q if, and only if, the value of $\{s^i\} \times Q^{-i}$ for player i attains or exceeds that of Q at every reachable position.*

Proof. Again, we argue for the two-player case. Given a strategy $s \in S^0$, we denote the value of $\{s\} \times Q^1$ for player 0 by χ_s , and the value of $Q^0 \times Q^1$ by χ .

It is easy to verify that the stated condition is sufficient. To prove that it is necessary, let us consider a strategy $s \in Q^0$ and a position p of minimal depth where $\chi_s(p) < \chi(p)$. (Recall that χ_s is defined only for positions reachable in $\{s\} \times Q^1$.) By definition of the value function, there exists a strategy $r \in Q^0(p)$ that witnesses the value of $Q^0 \times Q^1$ at p , having $\chi_r(p) = \chi(p)$. In particular, r agrees with s up to p . Since Q^0 allows shifting, it also contains the strategy $\sigma := s[p \leftarrow r]$. We claim that σ dominates s on Q^1 .

First, observe that for all positions p' incomparable to p , we have $\chi_\sigma(p') = \chi_s(p')$, whereas, for any p' comparable to p , we have $\chi_\sigma(p') = \chi_r(p')$. This already implies that the second condition of Lemma 8 is fulfilled: σ and s agree on p and $\chi_\sigma(p) = \chi_r(p) > \chi_s(p)$. The first condition of that lemma refers to the positions where σ and s split. We distinguish two cases. If $\chi_s(p) = 0$, there is only one split between s and σ that occurs immediately at p , where we know that $\chi_\sigma(p) > \chi_s(p)$. Otherwise, if $\chi_\sigma(p) = 1$, for any future position $p' \succeq p$ reached following σ , we have $\chi_\sigma(p') = 1$. Hence, whenever a split between σ and s occurs at such a position, we have $\chi_\sigma(p') \geq \chi_s(p')$. Accordingly, the first condition of Lemma 8 holds as well, allowing us to conclude that σ dominates s over Q^1 . \square

This characterisation implies that if a rectangular set $Q \subseteq S$ allows shifting, the set of strategies $s^i \in Q^i$ admissible with respect to Q also allows shifting. By transfinite induction it follows that the value characterisation is valid throughout all iteration stages of admissibility.

Corollary 10. *For every ordinal α , the stage Q_α allows shifting.*

4 One-step soundness

Next, we show that in any nonempty stage there exist admissible strategies and every eliminated strategy is dominated by one of those.

Theorem 11. *For every $r^i \in Q_\alpha^i \setminus Q_{\alpha+1}^i$ there exists a strategy $s^i \in Q_{\alpha+1}^i$ that dominates r^i on Q_α^{-i} .*

Proof. In a two-player setting, assume $r \in Q_\alpha^0$ is a strategy dominated on Q_α^1 . We construct an infinite chain $s_0, s_1, \dots \in Q_\alpha^0$, starting with $s_0 := r$, where every strategy s_ℓ is either dominated by its successor $s_{\ell+1}$, or equal to it. Intuitively, $s_{\ell+1}$ follows s_ℓ for the first ℓ steps and then, if the value at the current position can be improved, it switches to a strategy that achieves this. In terms of the

characterisation from Lemma 9, we thus ensure that $s_{\ell+1}$ attains the optimal value at any position p of depth at most ℓ that is reachable when the opponent plays in Q_α^1 , i.e., $\chi_{s_{\ell+1}}(p) = \chi(p)$, with our usual notation for the value of $\{s_{\ell+1}\} \times Q_\alpha^1$ and $Q_\alpha^0 \times Q_\alpha^1$ for player 0. Concretely, we set out with $s := s_\ell$ and look at each position p of depth ℓ reachable in $\{s_\ell\} \times Q_\alpha^1$. In case $\chi_{s_\ell}(p) < \chi(p)$, by definition of the value function, there exists a strategy $s' \in Q_\alpha^0$ with $\chi_{s'}(p) = \chi(p)$, and we update $s := s[p \leftarrow s']$. By performing these shifts wherever appropriate, we finally obtain a strategy s that yields the optimal value $\chi(p)$ at every position of depth ℓ reachable in $\{s_\ell\} \times Q_\alpha^1$ while preserving the – already optimal – value of $\{s_\ell\} \times Q_\alpha^1$ at any shorter play. Since Q_α^0 allows shifting, $s \in Q_\alpha^0$, and we set $s_{\ell+1} := s$.

Observe that in the above construction every position p at which we shift is minimal with the property that $\chi(p) > \chi_{s_\ell}(p)$. By the same argument as in the proof of necessity in Lemma 9, it follows that $s[p \leftarrow s']$ dominates s on Q^1 . Consequently, whenever shifts occurs between s_ℓ and $s_{\ell+1}$, the strategy $s_{\ell+1}$ dominates s_ℓ on Q_α^1 and, by transitivity, it dominates any previous strategy in the sequence. Also if, after some stage, the sequence stabilises, i.e., $s_\ell = s_k$ for every $k > \ell$, then s_ℓ is admissible with respect to $Q_\alpha^0 \times Q_\alpha^1$, according to Lemma 9. For the case in which the sequence does not stabilise, we consider its point-wise limit, i.e., the function σ assigning to position p of depth ℓ the value $\sigma(p) := s_{\ell+1}(p)$ (which further equals $s_k(p)$ for all $k > \ell$). Clearly, $\sigma \in S^0$. We can now verify that σ is admissible with respect to $Q_\alpha^0 \times Q_\alpha^1$, it belongs to Q_α^0 , and dominates r on Q^1 . \square

The above theorem implies that the elimination of dominated strategies will never evacuate all the strategies from a nonempty admissibility stage.

Corollary 12. *For every ordinal α , if $Q_\alpha \neq \emptyset$, then $Q_{\alpha+1} \neq \emptyset$.*

However, we remark that limit steps of the iteration may yield empty stages. As an example, consider a three-player game (an ω -centipede) with positions p_0, p_1, \dots all belonging to player 0 who can move from each p_k either to p_{k+1} or into a subgame Δ_k which stabilises after k stages with player 2 winning, player 1 pending, and player 0 losing. Further, let $p_0, p_1, \dots \notin W^0$ so that the strategy $p_k \mapsto p_{k+1}$ is eliminated in the first stage. Since each stage Q_{n+1}^0 should avoid the subgames $\Delta_0, \dots, \Delta_k$, player 0 will eliminate strategies at every finite stage ending up with $Q_\omega^0 = \emptyset$.

5 Progress and stabilisation

The example at the end of the previous section points out that the elimination of dominated strategies may require infinitely many iterations. In this section we provide an argument that mitigates this inconvenience, showing that the iteration actually never hangs: within a bounded number of stages, it always generates a definitive solution for some subgame. Towards this, we introduce a measure telling at which stage a subgame is solved from the viewpoint of player i , i.e., at which stage the iteration on that subgame closes.

Definition 13. Fix a game Γ . For a player i , the *closure ordinal* $\text{cl}^i(p)$ of a position p is the smallest ordinal α such that, for every stage $\beta \geq \alpha$, if p is reachable in Q_β^i , then $Q_\beta^i|_p = Q_\alpha^i|_p$. We write $\text{cl}(p)$ for the maximum $\text{cl}^i(p)$ over all players $i < n$.

If a position p is reachable in Q_∞ , then for $\alpha = \text{cl}^i(p)$ we have $Q_\infty^i|_p = Q_\alpha^i|_p$, i.e., from p onwards, any strategy that is α -admissible is iteratively admissible. Conversely, if a position p becomes unreachable at some stage Q_α , then $\text{cl}^i(p) \leq \alpha + 1$ for all players i . This can be interpreted as follows: when a play reaches p , it is apparent that the behaviour of some player is not $\alpha + 1$ -admissible. Consequently the hypothesis of common knowledge of rationality (incorporating admissibility) must be dropped and the iteration halts.

The notion of a partial solution implicit in the consideration of closures is intimately related to subgames. When instead of the original game Γ , we refer to a subgame Γ_p , we write $(Q_p)_\alpha^i$ for the corresponding iteration stages and $\text{cl}_p^i(\cdot)$ for the closure ordinals. The next lemma states that, for positions p reachable with admissible strategies, it does not matter whether we first iterate the elimination procedure on Γ and then restrict to the subgame Γ_p , or vice versa.

Lemma 14. *Let p be a position reachable in a stage Q_α . Then, with respect to the subgame Γ_p , we have $Q_\beta^i|_p = (Q_p)_\beta^i$ for every $\beta \leq \alpha$. If further $\text{cl}^i(p) \leq \alpha$, then $\text{cl}^i(p) = \text{cl}_p^i(p)$.*

It follows that at most two iteration stages may pass without discovering a subgame that closes for at least one player.

Lemma 15. *If, for an ordinal α there is no position p with $\alpha \leq \text{cl}^i(p) \leq \alpha + 1$, for some player i , then $Q_{\alpha+2} = Q_\infty$.*

Proof. We only sketch the argument. Notice that, if a subgame Γ_p does not close within two successive stages, the value χ of these stages does not change (if a position is found to be either winning or losing, it closes for the respective player within one more iteration step). Consequently, all subgames affected by the elimination of strategies in $Q_{\alpha+1} \setminus Q_{\alpha+0}$ are pending in Q_α and in $Q_{\alpha+1}$ and any eliminated strategy is losing in $Q_{\alpha+1}$. The strategies s , say of player 0, that are eliminated here are characterised as follows:

there exists a play $p_0, p_1, \dots \notin W^0$ following s that reaches a position $p_k \in T^0$ with $\chi^0(p_k) = 0$ such that no $p_\ell \in T^{-0}$ for $\ell > k$ has a successor $p' \neq p_{\ell+1}$ with $\chi^0(p') \geq 0$.

If we now consider the next elimination step, it turns out that the above criterion does not apply to any further strategy, because the triggering conditions depend only on values that do not change. Therefore, $Q_{\alpha+2} \setminus Q_{\alpha+1} = \emptyset$. \square

In the two-player case, whenever a game is closed for one player, it will also close for the other player within one step. With more players, the delay between closing stages can be arbitrary long. However, this delay can be bounded at least for some subgame.

Proposition 16. *Let p be a position reachable in a stage Q_α with $\text{cl}(p) \geq \alpha$. Then, there exists a position p' reachable from p such that $\alpha \leq \text{cl}(p') \leq \alpha + n$.*

Proof. By induction on n . First, apply Lemma 15 to the subgame Γ_{p_0} for $p_0 = p$, to find a reachable position p_1 which closes before $\alpha + 1$, for some player. Now, fix this player's strategy set and consider the subgame Γ_{p_1} as a game among $n - 1$ players to find a position $p_2 \succeq p_1$ that closes before $\alpha + 2$, and so on. \square

Applying Proposition 16 recursively for reachable subgames, we can conclude that the set of positions closed in any sequence of n iteration steps draws an antichain in the set of yet unresolved positions (with respect to the partial order of positions in the game tree).

Corollary 17. *Let $P(\alpha) := \{p \in T : \text{cl}(p) > \alpha\}$ be the set of positions that are not yet closed at stage α . Then, for every α , the difference $P(\alpha) \setminus P(\alpha + n)$ contains a maximal antichain in $P(\alpha)$.*

6 Infinite games on finite graphs

Finally, we discuss iterated admissibility on finite path-forming games with ω -regular objectives. Generalising the standard definition [10], we specify an n -player *graph game* as a finite directed graph $G = (V, E, (V^i)_{i < n})$ with a set of states V assigned to individual players according to the turn-partition $(V^i)_{i < n}$, and a set $E \in V \times V$ of edges. A play is an infinite path formed interactively on this graph starting from a fixed initial state and successively prolonged along edges chosen by the player to which the current state belongs. The objective of a player i is given by an ω -regular set of words $L_i \in V^\omega$ specifying the set of plays where he wins. We do not require the sets L_i to form a partition of V^ω . This model can be easily embedded into our extensive game setting by unravelling a graph G as a game tree $\mathcal{T}(G)$ with positions corresponding to finite paths through G starting from the initial position. We say that a position $p \in \mathcal{T}(G)$ *copies* a state $v \in G$, if p is an image of $v \in G$ under unravelling.

Parity games provide a canonical model of ω -regular graph games. An n -player parity game is represented by a graph over an initial segment V of the natural numbers, equipped with subsets $\Omega^i \subseteq V$, one for each player i , that specify the winning condition for a play v_0, v_1, \dots as follows: player i wins the play, if the least priority appearing infinitely often in the sequence v_0, v_1, \dots belongs to Ω^i . In the classical framework of two-player zero-sum games, it is well known that every regular game can be equivalently translated into a parity game. This remains true in the n -player non-zero sum setting, in the following sense: for every regular game $\mathcal{G} = (G, (L^i)_{i < n})$, we can construct a parity game $\widehat{\mathcal{G}} = (\widehat{G}, (\Omega^i)_{i < n})$ together with an isomorphism h from $\mathcal{T}(G)$ to $\mathcal{T}(\widehat{G})$ that naturally extends to strategies such that $(s^0, \dots, s^{n-1}) \in S$ is winning for a player i in \mathcal{G} if, and only if, $(h(s^0), \dots, h(s^{n-1}))$ is winning for i in $\widehat{\mathcal{G}}$; moreover, h preserves admissibility. Accordingly, the following results for parity games translate back to regular games.

The parity winning condition is prefix-independent: in a parity game \mathcal{G} , a play v_0, v_1, \dots is winning iff any suffix v_k, v_{k+1}, \dots is winning in the game \mathcal{G} starting from v_k . Without loss of generality, a subgame Γ_p in the extensive form of a parity game can be identified with the game $\Gamma \upharpoonright_p$ over the subtree of $\mathcal{T}(G)$ rooted at p . The value characterisation of admissible strategies then implies that the iterated elimination of dominated strategies in the extensive form of a parity game \mathcal{G} evolves in the same way on any two subgames rooted at copies of the same state in \mathcal{G} , provided they are simultaneously reachable.

Lemma 18. *Consider the extensive-form game obtained by unravelling a parity game \mathcal{G} , and let $p, p' \in \mathcal{T}(G)$ be two copies of the same state in \mathcal{G} . Then, for any admissibility stage Q_α in which p and p' are both reachable, we have $Q_\alpha \upharpoonright_p = Q_\alpha \upharpoonright_{p'}$.*

In particular, if some position $p \in \mathcal{T}(G)$ closes with $\text{cl}^i(p) = \alpha$, then every other copy $p' \preceq p$ of the same state also closes with $\text{cl}^i(p') = \alpha$. By Proposition 16, in every sequence of n iterations, at least one position closes. Since the original parity game has finitely many states, the iteration for admissible strategies on its extensive form will terminate after finitely many stages.

Proposition 19. *The closure ordinal of an n -player parity game with m states is at most $n \cdot m$.*

Strategies in a parity game can be represented as colourings of the ω -tree obtained by unravelling the finite game graph. The initial set of strategy profiles is recognisable by an ω -tree automata. Given an automaton that recognises a regular set Q of strategy profiles we can construct, with some routine, an automaton that recognises the profiles consisting of strategies admissible with respect to Q .

Proposition 20. *In a regular game, let $Q \subseteq S$ be a Cartesian product of ω -tree regular sets. Then, the set of strategies $s \in Q^i$ that are admissible with respect to Q is ω -tree regular.*

Accordingly, in any regular game, the stage Q_α^i is regular, for any finite α . By Proposition 19, the iteration of admissible strategies on a regular game over a finite graph terminates after a bounded number of stages. We can hence conclude that iteratively admissible strategies on such games can be recognised by finite-state tree automata.

Theorem 21. *For every regular game over a finite graph, the set of iteratively admissible strategy profiles is ω -tree regular.*

References

- [1] K. R. APT, *Order independence and rationalizability*, in TARK '05: Proc. Theoretical aspects of rationality and knowledge, 2005, pp. 22–38.
- [2] A. ARNOLD, A. VINCENT, AND I. WALUKIEWICZ, *Games for synthesis of controllers with partial observation*, Theoretical Comp. Science, 303 (2003), pp. 7–34.

- [3] A. BRANDENBURGER, A. FRIEDENBERG, AND H. J. KEISLER, *Admissibility in games*. To appear.
- [4] K. CHATTERJEE, *Two-player nonzero-sum omega-regular games.*, in CONCUR 2005 - Concurrency Theory, vol. 3653 of LNCS, Springer, 2005, pp. 413–427.
- [5] K. CHATTERJEE, T. A. HENZINGER, AND M. JURDZINSKI, *Games with secure equilibria.*, in Proc. LICS 2004, Logic in Computer Science, 2004, pp. 160–169.
- [6] K. CHATTERJEE, R. MAJUMDAR, AND M. JURDZINSKI, *On Nash equilibria in stochastic games.*, in Proc. Computer Science Logic – CSL 2004, vol. 3210 of LNCS, Springer, 2004, pp. 26–40.
- [7] D. GALE, *A theory of n-person games with perfect information*, Proceedings of the National Academy of Sciences, 39 (1953), pp. 495–501.
- [8] H. GIMBERT, *Jeux Positionnels*, PhD thesis, Université Paris 7, 2006.
- [9] R. J. GRETLEIN, *Dominance elimination procedures on finite alternative games*, International Journal of Game Theory, 12 (1983), pp. 107–113.
- [10] Y. GUREVICH AND L. HARRINGTON, *Trees, automata and games*, in Proc. of the 14th Annual ACM Symposium on Theory of Computing, STOC '82, ACM, 1982, pp. 60–65.
- [11] T. A. HENZINGER, *Games in system design and verification.*, in Proc. Theoretical Aspects of Rationality and Knowledge (TARK-2005), 2005.
- [12] J. HILLAS AND E. KOHLBERG, *Foundations of strategic equilibrium*, in Handbook of Game Theory with Economic Applications, R. Aumann and S. Hart, eds., vol. 3, Elsevier, 2002, ch. 42, pp. 1597–1663.
- [13] E. KOHLBERG AND J.-F. MERTENS, *On the strategic stability of equilibria*, Econometrica, 54 (1986), pp. 1003–37.
- [14] D. R. LUCE AND H. RAIFFA, *Games and Decisions*, Wiley, 1957.
- [15] G. J. MAILATH, *Do people play Nash equilibrium? Lessons from evolutionary game theory*, Journal of Economic Literature, 36 (1998), pp. 1347–1374.
- [16] L. M. MARX AND J. M. SWINKELS, *Order independence for iterated weak dominance*, Games and Economic Behavior, 31 (2000), pp. 324–329.
- [17] S. MOHALIK AND I. WALUKIEWICZ, *Distributed games*, in FSTTCS'03, vol. 2914 of LNCS, 2003, pp. 338–351.
- [18] L. P. ØSTERDAL, *Iterated weak dominance and subgame dominance*, Journal of Mathematical Economics, 41 (2005), pp. 637–645.
- [19] C. PAPADIMITRIOU, *Algorithms, games, and the internet*, in STOC '01: Proc. ACM Symposium on Theory of Computing, 2001, pp. 749–753.
- [20] W. THOMAS, *Infinite games and verification*, in Proc. CAV 2002 – Computer Aided Verification., vol. 2404 of LNCS, Springer, 2002, pp. 58–64.
- [21] M. UMMELS, *Rational behaviour and strategy construction in infinite multiplayer games*, Master's thesis, RWTH Aachen University, 2005.
- [22] I. WALUKIEWICZ, *A landscape with games in the background*, in Proc. Logic in Computer Science (LICS 2004), 2004, pp. 356–366.